

Heiko Leonhard

Real Estate Markets in the Blockchain Era

Real Estate Markets in the Blockchain Era



Dissertation zur Erlangung des Grades eines
Doktors der Wirtschaftswissenschaft

eingereicht an der Fakultät für Wirtschaftswissenschaften
der Universität Regensburg

vorgelegt von:
Heiko Leonhard

Berichterstatter: Prof. Dr. Wolfgang Schäfers
Prof. Dr. Bertram Steininger

Tag der Disputation: 09. Juli 2026

Heiko Leonhard

Real Estate Markets in the Blockchain Era

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
Doktor der Wirtschaftswissenschaft (Dr. rer. pol.)

Advisors:

Prof. Dr. Wolfgang Schäfers
Chair of Real Estate Management
IRE|BS International Real Estate Business School
University of Regensburg

Prof. Dr. Bertram Steininger
Department of Real Estate and Construction Management
KTH Royal Institute of Technology, Stockholm

University of Regensburg
IRE|BS International Real Estate Business School
Regensburg, Germany
Copyright © 2026 Heiko Leonhard.

Acknowledgments

While a doctoral journey is an individual challenge, mine would not have been possible without the support of many special people, to whom I am deeply grateful.

First of all, I would like to thank my primary supervisor, Prof. Dr. Wolfgang Schäfers, for giving me the initial opportunity to pursue my doctoral studies at the Chair of Real Estate Management at the IRE|BS Institute of the University of Regensburg. His constant mentorship, trust, and the opportunities he provided me with throughout these years will remain deeply appreciated. I would also like to thank my second advisor, Prof. Dr. Bertram Steininger, for his commitment to supervising my doctoral studies and for his valuable feedback and support throughout my academic journey and research.

My doctoral studies would not have been possible without the collaborative work of my co-authors, from whom I have learned so much. Therefore, I would like to thank Prof. Dr. Wolfgang Schäfers, who consistently guided me with his experience and strategic perspective, PD Dr. Maximilian Nagl for the inspiring collaboration at the beginning of my doctoral studies, and Prof. Dr. Gregor Dorfleitner for the fruitful research together. My special thanks go to Dr. Ralf Laschinger for his intellectual guidance during our many extensive discussions and for encouraging me to strive for the highest research standards.

Furthermore, I would like to thank the people at the institute, my colleagues, fellows, and ultimately, friends. With them, I was fortunate not only to share the day-to-day moments, which may have seemed insignificant at the time but, in retrospect, were perhaps the most meaningful in cheering one another, but also to attend unforgettable academic conferences around the world.

Finally, and above all, I would like to thank my father, Herbert, for inspiring and encouraging me to pursue this doctorate in the first place, as well as my mother, Maria-José, and my sister, Maria, for always being by my side. I am incredibly grateful to my girlfriend, Michelle, who has stood with me since the beginning of this journey and has shown endless patience, personal support, and understanding. Last but not least, my gratitude goes out to my friends for their patience and for giving me the much-needed contrast to my academically dominated everyday life.

Declaration on the Use of Artificial Intelligence (AI)

In the preparation of this dissertation, AI-assisted tools and large language models were used to improve spelling and grammar, and to support the linguistic optimization of my own wording. After using these tools, the content was independently reviewed and, where necessary, revised. I take full responsibility for the content of the submitted academic work. As of the time of writing, no official or specific guidelines of the Faculty of Business, Economics, and Real Estate of the University of Regensburg regarding the use of AI in dissertations have been adopted. Therefore, reference is made to the guidelines of the Faculty of Law of the University of Regensburg, dated May 7, 2025.

Contents

1	Introduction	1
1.1	Motivation, Background, and Framework	1
1.2	Course of Analysis and Research Questions	5
1.3	Research Contributions and Outputs	7
	References	9
2	Liquidity Mechanisms in Decentralized Finance	12
2.1	Introduction	13
2.2	Related literature and institutional details	14
2.3	Framework for execution cost and liquidity provision across mechanisms . . .	17
2.4	Data and summary statistics	20
2.5	Empirical analysis	27
2.6	Conclusion	40
	Appendix	42
	References	65
3	Tokenization, Transparency, and Time on Market	69
3.1	Introduction	70
3.2	Background on real estate tokenization and institutional marketplace details	71
3.3	Related literature and hypothesis development	74
3.4	Data	78
3.5	Methodology	84
3.6	Results	86
3.7	Conclusion	94
	Appendix	95
	References	99
4	Virtual Land in the Metaverse?	103
4.1	Introduction	105
4.2	Virtual land market	108
4.3	Data	110
4.4	Methodology	112
4.5	Results	113
4.6	Conclusion	133
	Appendix	135
	References	151
5	Conclusion	157
5.1	Executive Summary	157
5.2	Final Remarks and Outlook	162
	References	164

List of Tables

2.1	Trade statistics by market mechanism	23
2.2	Descriptive statistics	25
2.3	Transaction-level mechanism choice	28
2.4	Liquidity provision across market mechanisms	30
2.5	External liquidity provision in AMM pools	32
2.6	1-token price impact around the Terra collapse (AMM vs. P2P)	36
2.7	Bid-ask spreads around the Temporary Restraining Order (P2P vs. Buyback)	40
A.1	Overview of RWA tokenization platforms and their liquidity mechanisms . .	43
A.2	AMM event counts by pool quote token	45
A.3	Realistic P2P price signals order-book depth measures	48
A.4	Notation summary for the execution cost framework	50
A.5	Definition of all variables	57
A.6	Price signals and quasi-arbitrage	64
3.1	Summary statistics for property and market characteristics	81
3.2	Summary statistics for main regression variables	82
3.3	Determinants of Time on Market (ToM) – Sales	88
3.4	Determinants of remaining ToM after price revision	90
3.5	Revision strategy and selling premium	93
A.1	Determinants of Time on Market (ToM) – Removals	96
A.2	Determinants of Time on Market (ToM) – 90% sales-threshold robustness . .	97
A.3	Revision strategy and selling price (USD)	98
4.1	Transaction prices – Summary statistics	117
4.2	Realized returns per parcel – Summary statistics	119
4.3	Descriptive statistics of weekly returns	122
4.4	Determinants of correlation estimates	128
4.5	Comparison of portfolio performance metrics (annualized)	132
A.1	Global companies and XY-coordinates in Decentraland	135
A.2	Number of companies by business sectors	138
A.3	Number of companies by continent	138
A.4	Average monthly dollar transaction volume – Summary statistics	139
A.5	Average monthly number of transactions – Summary statistics	139
A.6	Description of variables	147
A.7	Performance metrics by year and asset (weekly basis)	148
A.8	Performance metrics by year and asset (annualized)	149
A.9	Correlation matrix of excess returns	150
A.10	Portfolio weights	150

List of Figures

1.1	Framework of the dissertation and real estate markets in the blockchain era .	4
2.1	RealT’s market evolution	21
2.2	Price signals and realized trade prices for one tokenized RWA (sell-side) . . .	22
2.3	Trading volume by market mechanism	24
2.4	Number of trades by market mechanism	24
2.5	AMM liquidity: Issuer vs external providers	31
2.6	P2P liquidity: Two-sided quote providers	33
2.7	AMM liquidity pool depth, P2P offer-book depth, and Buyback capacity . .	34
2.8	Event study of 1-token price impact around the Terra collapse	37
2.9	Event study of P2P bid–ask spreads around the Temporary Restraining Order	39
A.1	AMM mechanism	44
A.2	AMM: Levinswap web interface	45
A.3	Peer-to-peer (P2P) marketplace	46
A.4	P2P: YAM web interface	46
A.5	Buyback	49
A.6	Buy order: rational choice between AMM and P2P	53
A.7	Sell order: rational mechanism choice	54
3.1	Geographic distribution of RealT’s tokenized real estate	79
3.2	Listing price path: listing-level and spell-level (illustrative example)	83
3.3	Distribution of revision frequency and total price concessions per sold listing	84
3.4	Effect of price revision path on final selling premium	91
A.1	YAM peer-to-peer marketplace interface	95
4.1	Decentraland map and companies	110
4.2	Fashion district	111
4.3	Average transaction prices	114
4.4	Average transaction prices above \$10,000	115
4.5	Transaction prices	116
4.6	Realized returns per parcel 2019-2021	118
4.7	Realized returns per parcel 2022-2023	119
4.8	Number of transactions per parcel	120
4.9	Virtual land index vs REIT indices	121
4.10	Virtual land and REITs (Pearson)	124
4.11	Virtual land & FTSE Nareit Subset	125
4.12	Virtual land & FTSE Nareit Subset	125
A.1	Companies in Decentraland – Exemplary excerpt	137
A.2	Monthly transaction activity	139
A.3	Investor groups by average dollar transaction size	140
A.4	Virtual land & FTSE Nareit All Equity REITs Index	144

A.5	Virtual land & FTSE Nareit All Equity REITs Index	144
A.6	Virtual land & S&P Global REIT Index	145
A.7	Virtual land & S&P Global REIT Index	145
A.8	Virtual land & REITs (Kendall's τ)	146

Chapter 1

Introduction

1.1 Motivation, Background, and Framework

The organization of asset exchange through markets dates back to ancient civilizations, where economic transactions were documented on clay tablets. Thousands of years later, Smith (1776, p. 16) described exchange as a fundamental feature of human behavior, referring to the “certain propensity in human nature [...] to truck, barter, and exchange one thing for another.” Over time, history has shown that the way and the conditions under which such exchange takes place are changing profoundly and constantly, particularly with respect to how assets are represented, how transactions are executed and documented, and how market participants obtain and process information. A constant across the changing forms of exchange is that the exchange of assets involves both monetary and non-monetary costs and frictions, which markets both endogenously shape and are themselves shaped by.

These costs and frictions, and the way markets function, depend profoundly on the asset to be exchanged. Markets for heterogeneous, high-value, and location-specific assets such as real estate are characterized by infrequent trading, decentralized organization, and relatively high intermediation as the transaction process is complex and low-automated. This structure creates economically relevant frictions in search, matching, bargaining, and execution (Williams, 1995). As a result, such markets exhibit low liquidity, slow and intransparent price discovery, high transaction costs, and uncertainty regarding asset valuations and execution outcomes. In these respects, real estate and similar private asset markets differ fundamentally from public asset markets such as listed equities or standardized debt instruments and require special but often deficient attention (Goetzmann et al., 2021). At the same time, the global real estate market comprises an estimated aggregate market value of roughly \$400 trillion, exceeding that of all major investable asset classes combined and underscoring its importance for the global economy and investment landscape (Savills, 2025).

Because of the economic relevance of real estate and the persistence of these structural market frictions, there has been a recurring search for institutional, financial, and technological innovations that aim to enhance efficiency in real estate exchange, financing, and investment (Merton, 1995). In traditional real estate capital markets, institutional and financial concepts such as securitization, fund structures, and real estate investment trusts (REITs) have

demonstrated how illiquid real estate properties can be pooled, divided, standardized, and transformed into more liquid and easily tradable securities. As a result, indirect real estate investments became broadly accessible, allowing investors to reduce single-asset risk through asset pooling and portfolio diversification (Markowitz, 1952). The other side of the coin has shown that financialization and liquidity transformation through these financial products can introduce new structural risks and fragilities. These risks result from increased contagion due to broader capital-market integration, misaligned incentives in originate-to-distribute practices, and a liquidity mismatch between the underlying asset and the easier-to-trade structured claims (Purnanandam, 2011; Gorton and Metrick, 2012). The global financial crisis exposed the far-reaching consequences of these risks materializing at once. Another stream of more explicitly technological innovation has reshaped how direct, property-level real estate transactions are organized, intermediated, and processed (Baum, 2017). Economically, these innovations primarily affect information gathering, search and matching, intermediation, and pricing support, overall aiming to improve the efficiency of the exchange process. Examples include online listing and marketplace platforms as well as artificial-intelligence-based automation tools. Taken together, these institutional, financial, and technological innovations contribute to an ongoing financialization, digitalization, and also platformization of real estate markets.

This dissertation examines a more radical innovation with the potential to fundamentally alter the organization of real estate markets: blockchain technology. The overall aim of this cumulative dissertation is to interrogate conceptual propositions and prevailing market narratives around blockchain and real estate using real transaction data and rigorous empirical methods. Building on publicly available but technically complex blockchain data, this work is positioned at the forefront of an emerging empirical research agenda on blockchain-based real estate markets. The increasing market adoption, disruptive potential, and ongoing debate make it both timely and relevant for industry, academia, and policymakers. Unlike previous innovations that primarily transform the financial packaging of real estate exposure and its institutional embedding into broader capital markets or the property-level transaction process through digital platforms and tools, blockchain affects the financial, institutional, and technological dimensions at once. At a higher level of abstraction, blockchain may restructure asset and capital markets in general by moving modern finance and investment toward fully digital and programmable market infrastructure (International Monetary Fund, 2026). Conceptually, blockchain is a shared, decentralized, and immutable digital ledger in which transactions can be recorded and, through smart contracts, increasingly complex agreements and applications can be encoded and executed. While blockchains can be private and permissioned, much of the disruptive potential examined in this dissertation comes from the openness of public blockchain networks. In this sense, blockchain is more than Bitcoin (Nakamoto, 2008), a digital payment system designed to overcome the double-spending problem by the use of cryptography. With Ethereum (Buterin, 2014; Wood, 2014) and similar networks, blockchain has also enabled a decentralized financial market infrastructure referred to as decentralized finance (DeFi). These structural features matter economically because, once blockchain enables the representation of assets such as currencies, stocks, bonds, fund shares, or real estate as digital tokens, this representation changes how those assets can be recorded, verified, transferred, settled, and held. How blockchain alters existing market frictions remains an open empirical question. This dissertation addresses this question by examining blockchain-based real estate markets at the intersection of real estate economics, market microstructure research, and public blockchain data.

The digital representation of assets as blockchain-based tokens, referred to as tokenization, bridges DeFi with traditional asset and capital markets. In this context, the tokenization of physical and traditional financial assets is often referred to as real-world asset (RWA) tokenization, distinguishing it from purely digital blockchain-native assets such as cryptocurrencies and other virtual assets. While, in principle, a broad spectrum of assets can be tokenized, its economic implications for the asset and its exchange differ substantially (Cong et al., 2025b). For inherently liquid assets like currencies, stocks, or highly liquid commodities and bonds, the liquidity of their traditional markets primarily spills over to their blockchain-based outlets rather than being newly created there, as both economic intuition would suggest and first empirical evidence demonstrates (Cong et al., 2025a; Harvey et al., 2025). For tokenized illiquid assets at the opposite end of the asset liquidity spectrum, such as individual real estate properties, this market liquidity logic changes, as their traditional markets are illiquid in the first place. This raises the question of whether tokenization and blockchain-based trading mechanisms can create liquidity through fractionalization and the facilitation of secondary-market formation. This makes the tokenization of real estate a particularly revealing case for studying how familiar market structure norms could be altered. While fractionalization itself is not new, as seen in stocks or fund shares, tokenizing individual real estate properties is economically distinct because it extends divisibility and secondary-market tradability to assets that are natively non-standardized, lumpy, and illiquid. The resulting digital tokens represent economic claims on the underlying asset and become more accessible to a broader and global investor base with facilitated, transparent, and continuous online tradability.

In the real estate investment landscape, fractionalized tokenized exposures thereby form a new hybrid between direct property ownership and indirect pooled vehicles (Geltner et al., 2026; Liu and Chen, 2025). Unlike REITs or open-end funds, which provide portfolio exposure with limited discretion over individual assets, most real estate token offerings are currently linked to single properties, enabling more granular portfolio construction and governance (Swinkels, 2023; Kreppmeier et al., 2023). Relative to closed-end vehicles, tokenization allows for lower minimum investment thresholds and fewer effective lock-up frictions. Conceptually close to real estate crowdfunding in its single-asset, fractional exposure, tokenization embeds ownership and transfers on interoperable blockchain-based token standards, enabling transparent post-issuance trading and further utilization, such as collateralization (Schär, 2021), rather than confining claims within siloed crowdfunding portals with very limited secondary trading options (Lukkarinen and Schwenbacher, 2023). Tokenization may therefore move real estate closer to a novel and particular financial-market microstructure logic, making this perspective especially relevant for its analysis. The programmability and composability of tokenized real estate further enable innovative secondary-market designs and integration with other financial applications. At the same time, blockchain makes ownership changes, listings, and transactions more transparent, further shaping not only market architecture but also market outcomes. These characteristics create a new potential for liquidity transformation between hard-to-trade underlying assets and easy-to-trade claims, combining blockchain-native market mechanisms with familiar economic trade-offs, tensions, and fragilities known from banks and established investment vehicles (Diamond and Dybvig, 1983; Goldstein et al., 2017). Understanding how blockchain-based liquidity and trading mechanisms are designed, interact, and are organized for tokenized, formerly illiquid assets, such as real estate, remains under-researched and therefore forms the central research question of Paper 1, presented as Chapter 2 of this dissertation. Alongside these liquidity and

trading-related questions, the unusual transparency in blockchain-based real estate markets also allows classical real estate questions regarding listing behavior, time on market, and price adjustment to be revisited in a new digital environment. These transparency-related implications are investigated in Paper 2, presented as Chapter 3 of this dissertation.

Beyond the tokenization of RWAs, blockchain-based market forms also extend into a purely virtual realm. In particular, non-fungible tokens (NFTs) enable the ownership, exchange, and investment of natively virtual assets such as virtual art or virtual land (Goetzmann et al., 2026). As virtual environments, often referred to as the metaverse, have attracted increasing amounts of time, attention, and economic activity since the early internet (Castronova, 2002), the emergence of blockchain-based virtual economies and spatial assets within them becomes economically relevant in its own right. While virtual land lacks a physical counterpart, it exhibits several curious similarities to real estate, including scarcity, location-specificity, heterogeneity, transaction markets, and valuation shaped by surrounding activity, visibility, and user attention, while also serving as a virtual platform for interaction, commerce, and consumption (Goldberg et al., 2024; Goldberg et al., 2023). More broadly, related intangible asset markets such as internet domain names have already been analyzed from a real estate and urban economic perspective, suggesting that overlaps between virtual and physical spatial markets call for closer examination (Lindenthal, 2014; Lindenthal, 2018). In this sense, virtual land is not a direct real estate tokenization case, but rather a virtual extension of real-estate-like assets and market formation in virtual space. Paper 3, presented as Chapter 4 of this dissertation, therefore provides initial evidence on blockchain-based ownership, exchange, and spatial economics in purely virtual environments and their relation to physical real estate markets.




Figure 1.1: Framework of the dissertation and its papers within real estate markets in the blockchain era. Own illustration, based on Schulte et al. (2026, p. 729).

Figure 1.1 summarizes the conceptual framework of this dissertation by situating blockchain-based market forms within a broader taxonomy of real estate markets. The highlighted areas locate the three papers within this framework. Papers 1 and 2 focus on tokenized real estate, whereas Paper 3 examines virtual land as a blockchain-native market form. Across all papers, the common thread is how blockchain and tokenization create new forms of ownership, exchange, and market formation.

1.2 Course of Analysis and Research Questions

This section outlines the course of analysis of the dissertation and summarizes the research objectives and research questions of each paper. The overarching question of the three papers is how blockchain technology, especially via tokenization, reshapes real estate markets. The core market dimensions examined in the papers can be condensed into *secondary market design, liquidity, and trading*; *transparency, time on market, and listing behavior*; and the *purely virtual extension of real-estate-like assets* in the online world. Made feasible by around half a decade of transaction activity and data across blockchain-based real estate markets and platforms, this work manifests itself as the inception of an empirically focused research agenda.

Paper 1: Liquidity Mechanisms in Decentralized Finance: Design and Fragmentation in Illiquid Real-World Asset Markets

Paper 1 examines the largest tokenized real estate market of its time in terms of *secondary market design, liquidity, and trading*. The paper speaks to the broader class of illiquid RWAs, such as private equity, private credit, and other hard-to-trade assets, by analyzing real estate as an exemplary case. To do so, roughly 780,000 secondary-market transactions across three coexisting and common blockchain-based market mechanisms, namely automated market makers (AMMs), peer-to-peer (P2P) marketplaces, and redemption-like Buyback, are curated and analyzed to shed light on their mechanism design in isolation and the market fragmentation arising from their conjunction. Drawing from traditional market microstructure research, the core objective of the study is to examine the central aspiration of creating endogenous liquidity for formerly illiquid assets, particularly real estate, through the act of tokenization. The study is guided by the following research questions, which are derived from the three hypotheses developed in the paper:

- Do traders endogenously sort across execution mechanisms according to trade size and execution priorities, thereby fragmenting order flow and generating cross-mechanism price dispersion?
- How does trading activity shape the revenue intensity and profitability of external liquidity provision in AMM liquidity pools?
- How does execution stress affect the trading frictions arising from the coexistence of heterogeneous execution mechanisms?

Paper 2: Tokenization, Transparency, and Time on Market: Empirical Evidence from a Blockchain-Based Real Estate Marketplace

Paper 2 focuses on three key real estate market dimensions, namely *transparency*, *time on market*, and *listing behavior*, which are substantially altered through tokenization. Blockchain-based tokenization enables a digitally recorded environment in which not only lumpy real estate is divided into small-denominated tokens, increasing tradability, but listing, repricing, and execution behavior also become publicly observable. This market peculiarity motivates the study to revisit classic questions on overpricing and price adjustment in the Time on Market (ToM) literature that are far less observable in conventional property markets. To do so, the study extends the secondary market data from Paper 1 and reconstructs around 111,000 market listings, 72,000 sales, and detailed property data. The study is guided by the following overarching research questions, which summarize the central research focus and hypotheses developed in the paper:

- What are the determinants of ToM in a tokenized real estate marketplace?
- How do price revision strategies affect ToM and price outcomes when the full listing path is observable?

Paper 3: Virtual Land in the Metaverse? Exploring the Dynamic Correlation with Physical Real Estate

Lastly, Paper 3 explores the *purely virtual extension of real-estate-like assets*, referred to as virtual land, within a popular blockchain-based metaverse platform. Virtual land has attracted numerous companies and investors seeking to build virtual commercial space and explore new economic potentials. Given the underlying economic parallels between virtual land and physical real estate markets, this study examines the interconnected dynamics of both markets. Using a uniquely compiled dataset of 15,855 virtual land sales, a repeat sales index is constructed to analyze the risk-return patterns of this virtual spatial asset. Building on this index, the study employs coherence, regression, and portfolio simulations to investigate the dynamic correlation between virtual and physical real estate, its determinants, and possible portfolio implications. The study, therefore, addresses the following three research questions:

- Are virtual land market returns connected to physical land market returns?
- Does the connectedness vary for different investment horizons and time frames?
- Which drivers can explain the connectedness of virtual and physical land?

1.3 Research Contributions and Outputs

This section provides an overview of the three papers presented as dissertation chapters, including their authors, submission details, and conference presentations.

Paper 1: Liquidity Mechanisms in Decentralized Finance: Design and Fragmentation in Illiquid Real-World Asset Markets

Authors:

Ralf Laschinger, Heiko Leonhard, Gregor Dorfleitner, Wolfgang Schäfers

Submission Details:

Journal: *Management Science* (Special Issue on Digital Finance)

Status: Rejected with personal invitation to submit a revised new version (13 October 2025)
/ Resubmitted (12 February 2026) / Under review (5 May 2026)

Conference Presentations:

- 6th Annual Crypto and Blockchain Economics Research (CBER) Conference, New York City, USA (June 2026, accepted)
- American Finance Association (AFA) Annual Meeting, Philadelphia, USA (2026)*
- 32nd Pacific Rim Real Estate Society (PRRES) Annual Conference, Adelaide, Australia (2026)
- 7th Sydney Market Microstructure and Digital Finance Conference, Sydney, Australia (2025)*
- 38th Australasian Banking and Finance Conference, Sydney, Australia (2025)*
- MIT Symposium on Technology, AI, and Real Estate, Cambridge, USA (2025)
- 17th Real Estate Markets and Capital Markets Network (ReCapNet), Regensburg, Germany (2025)
- 31st Annual Meeting of the German Finance Association (DGF), Hagen, Germany (2025)
- World Finance Conference, Valletta, Malta (2025)*
- Munich Finance Day, Munich, Germany (2025)*
- 41st American Real Estate Society (ARES) Annual Conference, Tucson, USA (2025)
- Tokenomics Workshop on the Economics of Technology and Decentralization, Tokyo, Japan (2024)
- 30th European Real Estate Society (ERES) Annual Conference, Gdańsk, Poland (2024)

* Presented by a co-author.

Paper 2: Tokenization, Transparency, and Time on Market: Empirical Evidence from a Blockchain-Based Real Estate Marketplace

Authors:

Heiko Leonhard, Ralf Laschinger, Wolfgang Schäfers

Submission Details:

Journal: *Journal of Real Estate Research* (Special Issue on AI and Real Estate)

Status: Minor revision (23 April 2026)

Conference Presentations:

- 32nd European Real Estate Society (ERES) Annual Conference, Vienna, Austria (July 2026, accepted)
- 31st European Real Estate Society (ERES) Annual Conference, Athens, Greece (2025)

Paper 3: Virtual Land in the Metaverse? Exploring the Dynamic Correlation with Physical Real Estate

Authors:

Heiko Leonhard, Maximilian Nagl, Wolfgang Schäfers

Submission Details:

Journal: *Journal of European Real Estate Research* (Special Issue on Digital Technology and Innovation in Real Estate)

Status: Published (30 January 2025)

Conference Presentations:

- 40th American Real Estate Society (ARES) Annual Conference, Orlando, USA (2024)
- Doctoral Seminar of the Center of Finance, University of Regensburg, Regensburg, Germany (2023)
- 29th European Real Estate Society (ERES) Annual Conference, London, UK (2023)

References

- Ante L., Wazinski F.-P., and Saggu A. (2023). Digital real estate in the metaverse: An empirical analysis of retail investor motivations. *Finance Research Letters* 58, p. 104299.
- Barbon A. and Ranaldo A. (2026). On the quality of cryptocurrency markets: Centralized vs. decentralized exchanges. *Management Science. Articles in Advance*.
- Baum A. (2017). PropTech 3.0: The future of real estate. Report. Oxford: Saïd Business School, University of Oxford. URL: https://www.sbs.ox.ac.uk/sites/default/files/2019-03/FoRET-ReportSummary_0.pdf.
- Baum A. (2021). Tokenization – The future of real estate investment? *The Journal of Portfolio Management* 47.10, pp. 41–61.
- Bergkamp M., Sifat I., and Swinkels L. (2025). Market maturation and democratization effects of tokenized real estate. SSRN. Available at URL: <https://ssrn.com/abstract=5195172>.
- Buterin V. (2014). Ethereum: A next-generation smart contract and decentralized application platform. White Paper. URL: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- Case K. E. and Shiller R. J. (1987). Prices of single-family homes since 1970: New indexes for four cities. *New England Economic Review* September/October, pp. 45–56.
- Castronova E. (2002). On virtual economies. SSRN. Available at URL: <https://ssrn.com/abstract=338500>.
- Cong L. W., Landsman W. R., Rabetti D., Zhang C., and Zhao W. (2025a). Tokenized stocks. SSRN. Available at URL: <https://ssrn.com/abstract=5937314>.
- Cong L. W., Mayer S., and Rabetti D. (2025b). Tokenizing real-world assets. Working Paper.
- Diamond D. W. and Dybvig P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91.3, pp. 401–419.
- Dowling M. (2022). Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters* 44, p. 102097.
- Geltner D. M., Miller N. G., Van de Minne A., Eichholtz P., Lindenthal T., and Shen L. (2026). Commercial real estate analysis for investment, finance, and development. 4th ed. Routledge.
- Ghosh I., Alfaro-Cortés E., Gámez M., and García N. (2023). Do travel uncertainty and invasion rhetoric spur metaverse financial assets? Gauging the role of media influence. *Finance Research Letters* 51, p. 103434.
- Goetzmann W. N., Huang D., and Nozari M. (2026). Non-fungible tokens as investment. Working Paper 34837. National Bureau of Economic Research.
- Goetzmann W. N., Spaenjers C., and Van Nieuwerburgh S. (2021). Real and private-value assets. *The Review of Financial Studies* 34.8, pp. 3497–3526.
- Goldberg M., Schär F., and Kugler P. (2024). Land valuation in the metaverse: Location matters. *Journal of Economic Geography* 24.1, pp. 67–92.
- Goldberg M., Schär F., and Thürkauf D. (2023). Retailing and customer engagement in the metaverse: An empirical analysis. SSRN. Available at URL: <https://ssrn.com/abstract=4547100>.

- Goldstein I., Jiang H., and Ng D. T. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics* 126.3, pp. 592–613.
- Gorton G. and Metrick A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics* 104.3, pp. 425–451.
- Harvey C. R., Rabetti D., Lin C., and Zhang C. (2025). Tokenized gold. SSRN. Available at URL: <https://ssrn.com/abstract=5995434>.
- International Monetary Fund (2026). Tokenized finance. IMF Note 2026/001. International Monetary Fund. URL: <https://www.imf.org/en/Publications/IMF-Notes/Issues/2026/04/10/tokenized-finance-565468>.
- Knight J. R. (2002). Listing price, time on market, and ultimate selling price: Causes and effects of listing price changes. *Real Estate Economics* 30.2, pp. 213–237.
- Kreppmeier J., Laschinger R., Steininger B. I., and Dorfleitner G. (2023). Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *Journal of Banking and Finance* 154, p. 106940.
- Lindenthal T. (2014). Valuable words: The price dynamics of internet domain names. *Journal of the Association for Information Science and Technology* 65.5, pp. 869–881.
- Lindenthal T. (2018). Monocentric cyberspace: The primary market for internet domain names. *The Journal of Real Estate Finance and Economics* 57.1, pp. 152–166.
- Liu A. and Chen C. (2025). From real estate financialization to decentralization: A comparative review of REITs and blockchain-based tokenization. *Geoforum* 159, p. 104193.
- Lukkarinen A. and Schwienbacher A. (2023). Secondary market listings in equity crowdfunding: The missing link? *Research Policy* 52.1, p. 104648.
- Markowitz H. (1952). Portfolio selection. *The Journal of Finance* 7.1, pp. 77–91.
- Merlo A. and Ortalo-Magné F. (2004). Bargaining over residential real estate: Evidence from England. *Journal of Urban Economics* 56, pp. 192–216.
- Merton R. C. (1995). Financial innovation and the management and regulation of financial institutions. *Journal of Banking & Finance* 19.3–4, pp. 461–481.
- Nakamoto S. (2008). Bitcoin: A peer-to-peer electronic cash system. White Paper. URL: <https://bitcoin.org/bitcoin.pdf>.
- Ordano E., Meilich A., Jardi Y., and Araoz M. (2017). Decentraland: A blockchain-based virtual world. White Paper. URL: <https://decentraland.org/whitepaper.pdf>.
- Purnanandam A. (2011). Originate-to-distribute model and the subprime mortgage crisis. *The Review of Financial Studies* 24.6, pp. 1881–1915.
- Rosenberg N. (1996). Uncertainty and technological change. *Technology and Growth. Conference Series* 40. Boston, MA: Federal Reserve Bank of Boston, pp. 91–110.
- Savills (2025). How much is global real estate worth? Accessed 2 April 2026. URL: <https://impacts.savills.com/market-trends/how-much-is-global-real-estate-worth.html>.
- Schär F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review* 103.2, pp. 153–174.

- Schulte K.-W., Sebastian S., Allendorf G., Ropeter-Ahlers S.-E., and Leonhard H. (2026). Immobilieninvestition. Immobilienökonomie I: Betriebswirtschaftliche Grundlagen. Ed. by K.-W. Schulte, S. Bone-Winkel, and W. Schäfers. De Gruyter Oldenbourg, pp. 725–838.
- Smith A. (1776). An inquiry into the nature and causes of the wealth of nations. Vol. 1. London: W. Strahan and T. Cadell.
- Steininger B. I. (2023). Return–risk analysis of real estate tokens: An asset class of its own. *The Journal of Portfolio Management* 49.10, pp. 83–102.
- Swinkels L. (2023). Empirical evidence on the ownership and liquidity of real estate tokens. *Financial Innovation* 9, pp. 1–29.
- Vidal-Tomás D. (2022). The new crypto niche: NFTs, play-to-earn, and metaverse tokens. *Finance Research Letters* 47, p. 102742.
- Wheaton W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy* 98.6, pp. 1270–1292.
- Whitaker A. and Kräussl R. (2020). Fractional equity, blockchain, and the future of creative work. *Management Science* 66.10, pp. 4594–4611.
- Williams J. T. (1995). Pricing real assets with costly search. *The Review of Financial Studies* 8.1, pp. 55–90.
- Wood G. (2014). Ethereum: A secure decentralised generalised transaction ledger. Yellow Paper. URL: <https://ethereum.github.io/yellowpaper/paper.pdf>.

Chapter 2

Liquidity Mechanisms in Decentralized Finance: Design and Fragmentation in Illiquid Real-World Asset Markets

This research paper is co-authored with Ralf Laschinger (Institute for Financial Innovation and Technology, LMU Munich School of Management), Gregor Dorfleitner (Department of Finance, University of Regensburg), and Wolfgang Schäfers (IREBS, International Real Estate Business School, University of Regensburg), and is currently under review for the *Management Science* Virtual Special Issue on Digital Finance after a personal invitation to revise and submit a new version.

Abstract: The tokenization of real-world assets (RWAs) links traditional assets to a blockchain-based trading infrastructure, but whether sustainable liquidity can be created remains unclear. We study liquidity formation in a fully transparent tokenized market for illiquid RWAs where three market mechanisms coexist: automated market makers (AMMs), bilateral peer-to-peer (P2P) marketplaces, and issuer-provided Buyback. Guided by reduced-form frameworks of execution cost and liquidity provision, we analyze 777,038 secondary-market trades and contemporaneous market liquidity-states from 2019 to 2025. We find that fragmentation across mechanisms generates persistent trader sorting and price dispersion. Price signals and executed volume concentrate on P2P. AMMs provide immediacy but are structurally fragile: thin liquidity and low turnover in illiquid RWAs weaken liquidity-provider incentives, resulting in shallow, largely issuer-supported depth. During market stress, execution costs increase differentially across mechanisms, consistent with mechanism-specific liquidity constraints. Our findings show that tokenization alone does not ensure liquid secondary markets for illiquid assets and that liquidity outcomes become mechanism-specific under low turnover.

Keywords: Real-World Assets (RWA), Tokenization, Decentralized Finance (DeFi), Market Microstructure, Market Fragmentation

JEL: G10, G12, G14, O33, R30

2.1 Introduction

The tokenization of real-world assets (RWAs) changes not only how assets are represented, but how liquidity is created, allocated, and sustained. For traditionally illiquid, slow-moving, and heterogeneous assets, tokenization promises to create and extend secondary-market liquidity by enabling fractional ownership and continuous on-chain trading. At the same time, tokenization fragments potential liquidity across different execution technologies. In tokenized RWA markets, three fundamental liquidity mechanisms exist: automated market makers (AMMs), bilateral peer-to-peer (P2P) marketplaces, and issuer-provided buyback or redemption facilities (Buyback). These mechanisms differ sharply in pricing rules, depth formation, and execution certainty. While tokenization facilitates secondary markets in general, the central question is how the coexistence of heterogeneous liquidity mechanisms shapes trading and liquidity formation for tokenized illiquid RWAs.

In traditional finance, securitization has improved liquidity primarily for large, homogeneous, and frequently traded assets by aggregating demand and reducing participation frictions (O’Hara, 2003). For such assets, tokenization primarily reallocates existing liquidity by changing trading, settlement, and custody technology. By contrast, when tokenization is applied to private, heterogeneous assets such as real estate, private credit, or private equity, it attempts something more demanding: the endogenous creation of secondary-market liquidity for claims on slow-moving fundamentals (Case and Shiller, 1989; Williams, 1995; Whitaker and Kräussl, 2020; Goetzmann et al., 2021). In this setting, valuation opacity, search frictions, limited participation, and infrequent information updates remain binding constraints. While on-chain functionalities such as collateralization, borrowing, and transparent settlement can facilitate trading (Li et al., 2024), liquidity remains a property of market design rather than technology alone. Tokenization may therefore reallocate liquidity for liquid underlyings, but for illiquid RWAs it constitutes an attempt at liquidity transformation that raises fundamental coordination and incentive-compatibility challenges (Cong et al., 2025b).

Fragmentation is central to this problem. Theory shows that dispersed trading across venues can improve welfare when no single platform aggregates information efficiently, but also that segmentation can impair execution quality when deterministic pricing interacts with thin liquidity and limited participation (Malamud and Rostek, 2017; Chen and Duffie, 2021; Aoyagi and Ito, 2025). Empirical evidence from Decentralized Finance (DeFi) documents persistent cross-venue price dispersion and slow synchronization even for highly liquid crypto-native assets (Makarov and Schoar, 2020; Barbon and Ranaldo, 2026). In tokenized markets for illiquid RWAs, these forces are amplified. Trading frequency is low, participation is constrained by institutional frictions such as KYC-based whitelisting, and no single mechanism reliably aggregates information. In such environments, fragmentation does not merely redistribute order flow, but interacts with market design to determine whether liquidity can be financed and sustained at all.

This paper studies these issues in the context of RealT, one of the earliest and largest RWA tokenization platforms, focusing on real estate assets. The platform provides a rare empirical environment in which AMM, P2P, and Buyback mechanisms coexist under a uniform token design with full on-chain observability, providing systematic empirical evidence on the secondary-market trading of tokenized illiquid RWAs. This setting serves as a transparent laboratory for studying liquidity mechanisms that are commonly employed across

tokenized RWA platforms, allowing direct measurement of mechanism-specific execution frictions, trader sorting behavior, price dispersion, and liquidity provision. While not all tokenized RWA markets share the same configuration, many employ a combination of the liquidity mechanisms analyzed in this study, and observing their coexistence within a single transparent environment provides a clean setting to isolate how execution technology shapes liquidity in fragmented secondary markets (see Appendix A.1). To our knowledge, RealT is the only platform for tokenized illiquid RWAs with a track record of more than five years that simultaneously features all three liquidity mechanisms and fully publicly accessible blockchain transaction data, making it uniquely suited for comprehensive empirical analysis. We reconstruct transaction-level quotes, offers, pool states, and execution outcomes across all mechanisms, yielding a comprehensive transaction-level dataset. We discipline the empirical analysis with a strictly positive, reduced-form execution-cost framework that captures the core technological differences across mechanisms, complemented by a simple framework of liquidity provision.

The key takeaway of this study is therefore structural. In illiquid RWAs, tokenization alone does not ensure liquid secondary markets. Our results suggest that when different execution mechanisms coexist, traders sort endogenously, and liquidity formation is constrained by market design and incentive compatibility. Liquidity outcomes become mechanism-specific under low turnover. AMMs provide continuous quotes but fail to scale when trading activity is low, while P2P markets intermediate economically meaningful trades and concentrate price signals, but still resemble digitized search markets.

2.2 Related literature and institutional details

This section reviews the literature and institutional setting relevant to our analyses.

2.2.1 Literature

We organize the literature around the core tensions raised in the introduction: tokenization, DeFi microstructure, and market fragmentation.

Tokenization. Distributed ledger technology enables a decentralized and transparent infrastructure for asset representation: digital ownership transfer, automated execution of contracts, and real-time settlement (Buterin, 2014; Cong and He, 2019; Harvey et al., 2025). These innovations underpin the expansion of DeFi (Schär, 2021), yet the effectiveness of blockchain-based trading venues in fostering trading and liquidity formation for tokenized illiquid RWAs remains uncertain. The inherent characteristics of RWAs further complicate liquidity formation. Unlike cryptocurrencies, which are native to the blockchain, fungible, and actively traded across global markets, RWAs can be heterogeneous, exhibit low turnover, and depend on the integration of off-chain data such as valuations, legal records, and income flows. This dependency introduces additional frictions through whitelisting and price oracles (Harvey and Rabetti, 2024; Cong et al., 2025c). However, first empirical evidence from tokenized markets indicates that these markets struggle with thin trading volumes, while showing early signs of maturation (Swinkels, 2023; Bergkamp et al., 2025). Recent

evidence on tokenized equities shows that blockchain-based markets can incorporate value-relevant information (Cong et al., 2025a), while tokenization may not uniformly create liquidity but instead reallocates or transforms liquidity depending on asset characteristics (Cong et al., 2025b). Despite a growing literature on tokenized assets, empirical research on secondary-market liquidity for RWAs remains sparse. Most studies focus on primary issuance mechanisms, such as Security Token Offerings (STOs), and their role in lowering investment barriers through fractional ownership (Lambert et al., 2022; Kreppmeier et al., 2023). These studies leave limited evidence on how liquidity forms and how execution outcomes differ across trading mechanisms once assets begin trading among investors.

DeFi microstructure. Conventional centralized exchanges aggregate liquidity through limit order books (Kyle, 1985; O’Hara, 2003), complemented by dealer-style venues like over-the-counter (OTC) systems for less liquid markets (Duffie et al., 2005). Blockchain-based decentralized exchanges implement analogous functions through smart contracts. AMMs are the dominant DeFi exchange design. They use algorithmic pricing to provide continuous on-chain quotes but can sustain persistent price deviations under trading frictions and expose traders to front-running attacks (Malinova and Park, 2017; Park, 2023; Barbon and Ranaldo, 2026; Capponi and Jia, 2025). A complementary class of decentralized venues are on-chain P2P protocols that resemble OTC execution. Existing studies primarily examine liquidity determinants in cryptocurrency markets, both on centralized (Brauneis et al., 2021) and decentralized exchanges (Zhu et al., 2025). However, little research has addressed how these liquidity mechanisms operate and interact in tokenized RWA markets. The coexistence of multiple liquidity mechanisms introduces structural complexity that can distort pricing and execution outcomes, motivating further investigation as more traditional assets move on-chain into an increasingly complex DeFi ecosystem.

Market fragmentation. Building on canonical models of coexisting trading platforms (Pagano, 1989; Malamud and Rostek, 2017; Chen and Duffie, 2021; Aoyagi and Ito, 2025), we examine how fragmentation across venues shapes trading and liquidity formation. These models show that when investors optimally split orders, a finite degree of fragmentation can improve aggregate trading outcomes even as prices on individual venues become less informative. At the same time, excessive segmentation diffuses information and increases arbitrage exposure (Makarov and Schoar, 2020). Extending these insights, Rostek and Yoon (2025) show in a decentralized multi-asset setting that product design and cross-security inference interact with price impact, so market quality may still improve under limited fragmentation even as individual-venue depth declines. Prior research has largely focused on pricing and information aggregation in centralized trading mechanisms (Kyle, 1985; Hasbrouck, 1995; O’Hara, 2003) or decentralized mechanisms studied in isolation (Capponi et al., 2026). How price information transmits across fundamentally different market designs remains largely untested. In particular, there is limited evidence on whether deterministic pricing, bilateral search, and appraisal-based quoting can jointly facilitate trading and support liquidity formation. Markets for tokenized illiquid RWAs provide a transparent setting in which these theoretical tensions can be tested empirically.

2.2.2 Liquidity mechanisms and market design

Three distinct liquidity mechanisms coexist in our setting, operating under fundamentally different designs: AMMs, P2P, and Buyback.

AMMs. Protocols such as Uniswap V2 (Adams et al., 2020) provide continuous two-sided liquidity at deterministic prices set by the constant-product market-making (CPMM) rule (see Appendix Figure A.1). Pool quotes are mechanically determined by reserve ratios, so price signals are fully embedded in the CPMM curve and adjust deterministically to order flow. Execution is immediate, but slippage increases convexly with trade size. Liquidity providers (LPs) earn fees proportional to volume and may incur impermanent loss when relative token prices move asymmetrically (Capponi and Jia, 2021). Our setting includes three constant-product AMMs: Uniswap V1, Uniswap V2, and, primarily, Levinswap, a Uniswap V2 fork deployed on the Gnosis blockchain (Gnosis), a compatible blockchain connected to Ethereum mainnet (Ethereum) (see Appendix Figure A.2 for an interface example). For each RWA token, the issuer creates and seeds a single liquidity pool, predominantly quoted against the USDC stablecoin, to enable immediate trading and attract external liquidity providers. Following a successful STO, pools are initialized with fixed issuer-sponsored liquidity. As the underlying RWA tokens generate cash flows, mainly rental income, LPs receive the income stream associated with the tokens they contribute, plus a pro-rata share of the cash flow tied to the issuer-sponsored position, intended to incentivize participation beyond trading fees.

P2P. Marketplaces such as YAM and Swapcat support bilateral, posted-offer trading with user-defined prices and quantities (see Appendix Figure A.3 for an illustrative example). They operate as decentralized posted-offer markets: participants list non-negotiable, take-it-or-leave-it buy or sell offers, creating observable bid and ask signals without automated matching. Execution is therefore stochastic and depends on search frictions, posted depth, and deviations from appraisal-based or AMM-based prices. This structure resembles decentralized OTC trading: it offers flexible pricing, transparent quoting, and support for larger or less time-sensitive trades. The buy side is typically thin and often hidden, creating the structural asymmetry that later shapes our empirical analyses. The corresponding P2P price signals are posted bid and ask quotes, closely resembling OTC-style indicative offers that reveal sellers' reservation values and the sparse, partially hidden buy-side demand. For details on the offer book interface, see Figure A.4 in Appendix.

Buyback. The issuer offers investors the option to sell RWA tokens back at an appraisal-based valuation, resembling a fund-style redemption mechanism (see Figure A.5). Pricing is stable but infrequently updated, execution requires a deterministic waiting period, and liquidity is subject to maximum caps, for example \$2,000 per wallet per week in our setting. While this mechanism guarantees execution at a known price, it offers limited flexibility and weak real-time price signals, making it most suitable for larger redemptions, simple execution needs, or traders prioritizing certainty over market-improved pricing. The Buyback price signal is the appraisal-anchored valuation, which is explicitly tied to the externally certified asset appraisal divided by the outstanding token supply, and serves as a reference point for secondary-market pricing. Importantly, access to all three secondary-market mechanisms is permissioned. Wallets and RWA tokens must be whitelisted through the issuer platform's KYC process, and only whitelisted wallets can trade or provide liquidity in the corresponding RWA tokens. For more institutional details, see Appendix A.1.1.

2.3 Framework for execution cost and liquidity provision across mechanisms

This section focuses on the key economic implications that discipline the empirical analyses, while the Appendix A.2 provides a detailed exposition of the framework.

2.3.1 Execution cost framework

Trader sorting.

Each trader i , who wishes to place a buy or sell order for a quantity q of a certain RWA token is exposed to three sources of execution cost: (i) proportional fees, (ii) price impact (slippage), and (iii) execution delay. All mechanisms share a common reference price P .¹ Since every total-cost expression contains the factor $qP > 0$, we normalize costs by qP throughout and work with mechanism-specific normalized execution costs.

AMMs. Let L_A denote AMM liquidity (RWA token units). Under the constant-product rule, the post-trade price after a buy or sell of size q is

$$P_{\text{new}} = \begin{cases} P \left(\frac{L_A}{L_A - q} \right)^2, & \text{for a buy,} \\ P \left(\frac{L_A}{L_A + q} \right)^2, & \text{for a sell.} \end{cases} \quad (2.1)$$

Under the assumption that the trader's average execution-price impact is approximated by one half of the marginal price change from P to P_{new} , the normalized AMM execution costs are therefore

$$\begin{aligned} c_A^{\text{buy}}(q) &= F_A + \frac{1}{2} \left[\left(\frac{L_A}{L_A - q} \right)^2 - 1 \right], \\ c_A^{\text{sell}}(q) &= F_A + \frac{1}{2} \left[1 - \left(\frac{L_A}{L_A + q} \right)^2 \right]. \end{aligned} \quad (2.2)$$

P2P with depth ($L_P > 0$). With posted offer depth L_P , execution incurs sub-quadratic price impact PI^+ (Almgren and Chriss, 2001):

$$c_P^{(+)}(q; L_P) = F_P + \frac{1}{2} PI^+ \left(\frac{q}{L_P} \right)^\delta, \quad \delta \in (0, 1), \quad (2.3)$$

so impact scales as $q^{1+\delta}$ ("walking the book") and decreases with available depth L_P . All execution-cost expressions and routing conditions for the case of zero P2P depth ($L_P = 0$) are provided in Appendix A.2.

Buyback. Buyback execution occurs at a fixed relative discount F_B and deterministic waiting time τ_B . Delay enters linearly via the trader's individual urgency preference λ :

$$c_B = F_B + \lambda \tau_B. \quad (2.4)$$

¹Note that price dispersion across mechanisms is captured by an additional parameter for transaction costs referring to the AMM price.

Routing conditions. Traders choose the mechanism with the lowest normalized cost.

Sell orders. AMM is preferred to Buyback when

$$\lambda < \frac{F_A - F_B + \frac{1}{2} - \frac{1}{2}(\frac{L_A}{L_A+q})^2}{\tau_B}. \quad (2.5)$$

With depth ($L_P > 0$), P2P is preferred to Buyback when

$$\lambda > \frac{F_P - \frac{1}{2}PI^+(q/L_P)^\delta - F_B}{\tau_B}. \quad (2.6)$$

Finally, P2P dominates AMM for $L_P > 0$,

$$1 - \left(\frac{L_A}{L_A+q}\right)^2 - PI^+\left(\frac{q}{L_P}\right)^\delta > 2F_P - 2F_A. \quad (2.7)$$

Formal derivations of these comparative statics and buy order routing conditions are provided in Appendix A.2.

2.3.2 Liquidity provision framework

P2P markets can in principle support voluntary market making, but when trading frequency is low, incentives to supply continuous bid and ask liquidity are weak, mirroring traditional OTC markets. By contrast, AMMs provide continuous quotes mechanically, but the depth of those quotes depends on liquidity supplied by LPs. Thus, LP incentives for external liquidity provision are economically central. While AMMs are well studied for crypto-native assets, their viability for tokenized illiquid RWAs remains largely unexplored. We therefore focus on AMM liquidity provision and develop a simple framework that links external liquidity to turnover, fee revenue, and issuer-sponsored incentives. In each liquidity pool, the token issuer contributes L_R units of RWA tokens, paired with an equal-value amount of quote tokens (here a stablecoin) at reference price P . The issuer further subsidizes liquidity provision by transferring the cash-flow yield associated with L_R to external LPs, allocated pro rata by liquidity shares. An external LP then chooses additional liquidity L in RWA tokens, together with the matched quote amount, so that total RWA-side liquidity equals $L_R + L$. Over the investment horizon, the pool experiences a (ex ante unknown) total buy q_b and sell q_s volume of RWA tokens. We summarize the *net order imbalance* as

$$\Delta q := q_b - q_s, \quad (2.8)$$

which is positive (negative) when buy (sell) orders dominate. We define *total trading activity* as

$$Q := q_b + q_s, \quad (2.9)$$

capturing overall turnover in the pool over the investment horizon. Both Δq and Q must be estimated by the external LP when choosing how much liquidity to supply.

The external LP's profit over the horizon has four components: (i) own and prorated yield on the RWA tokens supplied to the pool; (ii) fee income from total trading volume; (iii)

valuation changes arising from the constant-product pricing rule²; and (iv) the opportunity cost of capital required to finance the liquidity provision. Here, μ is the *yield* earned by holding the RWA tokens, and r is the *LP's cost of capital*. The term $a = \mu - 2r$ captures the net valuation effect of holding base tokens rather than investment alternatives. The linear coefficient b incorporates (i) the valuation effect of issuer seed liquidity, $2L_R(\mu - r)$, (ii) expected fee income QF_A from trading activity Q at AMM fee rate F_A , and (iii) the valuation effect of expected net order flow Δq . The constant term $c = \mu L_R^2$ reflects the cash-flow value of issuer-supplied liquidity.

Collecting terms and normalizing by the reference price P , expected profit for an external LP takes the form

$$\Pi(L) = \frac{aL^2 + bL + c}{L_R + L}, \quad (2.10)$$

where the coefficients collect distinct economic effects:

$$\begin{aligned} a &= \mu - 2r, \\ b &= 2L_R(\mu - r) + QF_A + \Delta q\left(2 - \frac{\mu}{2}\right), \\ c &= \mu L_R^2. \end{aligned} \quad (2.11)$$

For convenience, we define a *revenue intensity* term

$$\gamma := QF_A + \Delta q\left(2 - \frac{\mu}{2}\right), \quad (2.12)$$

which bundles the turnover-driven fee component and the valuation effects associated with directional order flow. Under the empirically relevant condition $2r > \mu$, the profit function $\Pi(L)$ is concave in L .³ If the revenue-intensity satisfies $\gamma > 0$, the optimal external liquidity is given by an expression which tends to be negative for smaller values of γ . More importantly, the full set of liquidity choices for which the external LP earns weakly positive profits for all $L \in [0, L']$ can be characterized by

$$\begin{aligned} L' &= \frac{2L_R(\mu - r) + \gamma}{2(2r - \mu)} + \\ &\quad \frac{\sqrt{4L_R^2 r^2 + 4L_R(\mu - r)\gamma + \gamma^2}}{2(2r - \mu)}. \end{aligned} \quad (2.13)$$

The expression shows that AMM liquidity increases with the issuer's seed liquidity L_R and with expected revenue intensity γ . Because $\partial L' / \partial L_R > 0$, issuer-sponsored liquidity expands the set of $(Q, \Delta q)$ realizations under which external LPs can profitably supply liquidity. We provide full derivations of the equations in Appendix A.2.1.

²Note that we do not use the usual “impermanent loss” concept, as its standard benchmark is holding the same quantities of the base and quote tokens outside of an AMM (i.e., a passive two-asset hold). This behavior may be typical of a crypto investor with a diversified token portfolio and is useful as a relative performance measure of providing liquidity versus holding that 50/50 portfolio. In our setting, however, the economically relevant decision is whether to commit fresh capital to acquire and tie up the inventories required to provide liquidity, or instead to hold the RWA token (or invest funds elsewhere). We therefore framework liquidity provision as an absolute profit problem and explicitly incorporate the opportunity/financing cost of capital at rate r . This automatically limits the liquidity that can be provided due to capital costs; in an approach based on impermanent loss, however, there are no such limits.

³This condition requires that the investor's capital costs exceed the AMM's passive return. This is empirically plausible in real-world settings, as otherwise a form of unbounded arbitrage would emerge (raising capital at rate r and earning $\frac{\mu}{2}$).

2.3.3 Hypotheses

Based on the theoretical frameworks, the following hypotheses will guide our empirical analyses:

Hypothesis 1 (Sorting, fragmentation, and persistent price dispersion) *Traders endogenously sort across execution mechanisms according to trade size and execution priorities, fragmenting order flow and generating cross-mechanism price dispersion.*

Hypothesis 2 (Turnover-driven AMM liquidity fragility) *Low trading activity reduces the revenue intensity of AMM pools and shrinks the profitability for external LPs. When turnover is insufficient to cover opportunity costs, AMM liquidity remains shallow.*

Hypothesis 3 (State-dependent amplification under stress) *Under execution stress, the coexistence of heterogeneous execution mechanisms amplifies trading frictions by intensifying order-flow fragmentation, leading to systematically larger divergences in execution outcomes.*

2.4 Data and summary statistics

This section describes the data and presents summary statistics used in the empirical analyses.


2.4.1 Platform

We study RealT, a tokenization platform that issues and trades RWAs on public blockchains. RealT has operated since 2019 and, by October 2025, had tokenized approximately \$163 million worth of underlying assets, primarily rental residential real estate in the United States and Latin America. Its long track record and advanced ecosystem featuring multiple secondary-market mechanisms make it a unique data-rich laboratory for studying secondary markets for tokenized illiquid RWAs. More importantly, because many other tokenization platforms rely on a combination of these same mechanism categories, our study is naturally transferable beyond the RealT setting (see Appendix A.1).

All issuances and secondary-market trades occur on Ethereum and Gnosis and are executed via smart contracts, ensuring full on-chain traceability and allowing us to reconstruct complete transaction-level panels from November 2019 to October 2025, comprising 777,038 trades, \$41.2 million in traded volume, and 9,202 unique wallet addresses. Figure 2.1 summarizes the platform’s aggregate market evolution.

We enrich each secondary-market trade with contemporaneous, liquidity-relevant market states by reconstructing pool-level AMM reserves and liquidity-provider positions, on-chain P2P offer books with standing bid and ask inventories, and Buyback flows. This yields mechanism-specific panels matched directly to individual trades.

Figure 2.1: RealT’s market evolution



Note: This figure plots three global series for the RealT token universe (million USD). *Total market capitalization* is computed each day as the sum across tokens of total token supply times a 7-day moving average of the token’s daily secondary-market volume-weighted average trading prices. *Cumulative primary market issuance volume* is the sum across tokens of the issuance value at STO. *Cumulative secondary market trading volume* is the cumulative sum of secondary-market trading volume. Values are forward-filled over time for missing observation days.


At the token level, we merge these data with detailed token- and asset-level fundamentals for 752 RWA tokens, combining time-invariant attributes (e.g., STO terms and property characteristics) with time-varying fundamentals (e.g., rental income and appraisal values). Mechanism overlap is pervasive: 97% of tokens trade on both P2P and AMMs, and 73% trade on all three mechanisms. We aggregate the resulting components into token- and pool-level panels at multiple frequencies for the empirical analyses. Appendix A.4 provides full details on database construction, validation procedures, and reproducibility.

2.4.2 Price signals and trades

Figure 2.2 illustrates, for a single representative tokenized RWA and its tokens, the core implications of our theoretical framework. Each liquidity mechanism generates distinct price signals and trade prices because execution costs, delay, and depth differ systematically across mechanisms. Even for an appraisal-anchored asset, price signals display substantial dispersion, and realized trade prices fluctuate widely over time.

Buyback prices remain flat, reflecting infrequent revaluations. In contrast, AMM price signals vary because quotes are mechanically tied to liquidity-pool reserves: even small order imbalances push pool prices away from appraisal values or contemporaneous P2P offer prices, creating temporary mispricing and quasi-arbitrage opportunities (see Appendix A.5). With both offer-sides present, P2P price signals display a persistent bid–ask spread, with asks reflecting seller-posted offers and bids representing posted demand. On the buy-side, filled P2P bids are rare: P2P buy-side liquidity is mostly hidden rather than posted.

Figure 2.2: Price signals and realized trade prices for one tokenized RWA (sell-side)



Note: This figure plots daily USD price signals and realized trade prices for one exemplary tokenized RWA on the *sell-side* in our sample. Observation counts (in parentheses) report the number of available observations underlying each series and sum to 20,497 price-signal observations and 7,024 realized trades (trade prices). The Buyback quotes based on appraisal values (a flat line over the sample period because no re-appraisal occurred; $N = 7,024$), AMM quotes implied by pool reserves (identical for bid and ask side; $N = 7,024$), P2P bid ($N = 968$) and ask ($N = 5,481$) signals computed as daily median offer prices from reconstructed offer books, and realized trade prices ($N = 7,024$). Parentheses denote P2P ask signals, which are not executable on the sell-side.

In the full sample, Table 2.1 documents large and systematic differences in trading activity across mechanisms, with highly skewed intraday trading. Overall, individual trade sizes are typically very small (mean 0.97 tokens; median 0.03), with economic volume concentrated in a narrow right tail: the median trade volume is \$1.30, while the maximum exceeds \$89,000. AMMs handle approximately 71% of all trades by count, P2P 27%, and Buyback only 2%, with nearly all activity settling on the low-cost Gnosis (99%). This aggregate skewness maps cleanly into sharp cross-mechanism differences. AMM trades are extremely numerous but very small (median 0.016 tokens; mean volume \$6.26), consistent with their role in servicing immediacy-driven, low-value trades. P2P trades are roughly an order of magnitude larger (median 1 token; mean volume \$166.87) and account for the majority of economic volume. Buyback trades are infrequent but involve the largest average trade sizes. P2P trades clear at the highest average prices (\$56.42), while AMM prices exhibit the greatest dispersion, reflecting thin pools and episodic mispricing. Cumulatively, P2P accounts for \$34.7 million (84%) in total traded volume—an order of magnitude more than either AMMs (8%) or Buyback (7%).


Table 2.1: Trade statistics by market mechanism

Statistic	N	Mean	St. Dev.	Min	10%	Median	90%	Max
<i>Panel A: AMM</i>								
Trade size (RWA token)	555,169	0.099	0.668	0.000	0.009	0.016	0.100	300.000
Trade volume (USD)	555,169	6.261	43.915	0.000	0.488	0.824	6.387	3,782
Total fees (USD)	555,169	0.048	0.755	0.000	0.002	0.003	0.022	117.244
Trade price (USD/RWA token)	555,169	54.551	15.752	0.184	47.189	52.835	61.287	2,394
<i>Panel B: P2P</i>								
Trade size (RWA token)	207,885	3.068	14.652	0.000	0.039	1.000	6.000	1,771
Trade volume (USD)	207,885	166.870	788.185	0.000	2.000	51.610	333.960	89,970
Total fees (USD)	207,885	0.192	2.827	0.000	0.000	0.000	0.003	148.697
Trade price (USD/RWA token)	207,885	56.422	16.444	1.000	49.000	52.555	65.000	617.000
<i>Panel C: Buyback</i>								
Trade size (RWA token)	13,984	4.134	6.431	0.000	0.157	2.000	10.000	40.000
Trade volume (USD)	13,984	219.434	341.807	0.000	8.041	100.600	520.140	2,000
Total fees (USD)	13,984	7.690	11.478	0.000	0.256	3.057	19.218	118.082
Trade price (USD/RWA token)	13,984	53.189	13.000	44.190	49.560	50.715	54.950	185.800
<i>Panel D: Total Market</i>								
Trade size (RWA token)	777,038	0.966	7.771	0.000	0.009	0.029	2.000	1,771
Trade volume (USD)	777,038	53.066	418.581	0.000	0.491	1.305	100.000	89,970
Total fees (USD)	777,038	0.224	2.437	0.000	0.000	0.003	0.021	148.697
Trade price (USD/RWA token)	777,038	55.027	15.918	0.184	47.818	52.728	62.437	2,394
<i>Cumulative trading volumes (USD, full sample)</i>								
Total trading volume (AMM, USD)								3,476,186
Total trading volume (P2P, USD)								34,689,710
Total trading volume (Buyback, USD)								3,068,572
Total trading volume (All, USD)								41,234,467

Note: Panels A–C report summary statistics for *Trade size (RWA token)*, *Trade volume (USD)*, *Total fees (USD)*, and *Trade price (USD/RWA token)* by mechanism (AMM, P2P, Buyback). Panel D reports the same variables for the full sample and additionally the *Cumulative trading volumes (USD)* by mechanism. We report the 10th percentile and 90th percentile to highlight distributional skewness. All statistics exclude missing observations.


In aggregate form, these patterns scale into clear market-level trends. As shown in Figures 2.3 and 2.4, AMMs dominate by *number of trades* but not by *trading volume*: they intermediate many small, immediacy-driven trades, whereas P2P venues execute far fewer but substantially larger trades. This separation reflects the underlying execution technologies. Beginning in late 2021, lower transaction costs on Gnosis and maturing P2P infrastructure led to a pronounced shift of volume toward P2P venues. By early 2022, P2P had become the dominant mechanism by traded USD volume, while AMMs persisted mainly for small, time-sensitive flows. When not suspended, the Buyback remained a stable but infrequent outlet for large redemptions.

Figure 2.3: Trading volume by market mechanism



Note: This figure presents the monthly trading volume in million USD categorized by market mechanisms.

Figure 2.4: Number of trades by market mechanism



Note: This figure presents the monthly number of trades categorized by market mechanisms.

2.4.3 Descriptive statistics

Table 2.2 presents descriptive statistics at the intraday (transaction-level) and daily aggregation levels for the key variables used throughout our analyses and subsample constructions, including trading outcomes, mechanism-specific liquidity-states and measures, and market control variables. Overall, our data set covers 752 distinct tokenized RWAs, totaling approximately 3 million tokens available for trading.⁴

Table 2.2: Descriptive statistics

Statistic	N	Mean	St. Dev.	Min	10%	Median	90%	Max
Trades-intraday perspective								
Trade size (RWA token)	777,038	0.97	7.77	0.00	0.01	0.03	2.00	1,771
Trade volume (USD)	777,038	53.07	418.58	0.00	0.49	1.30	100.00	89,970
RWA token buy	777,038	0.38	0.49	0.00	0.00	0.00	1.00	1.00
Ethereum blockchain	777,038	0.01	0.10	0.00	0.00	0.00	0.00	1.00
Price dispersion-intraday perspective								
Price signal AMM (USD)	776,532	62.44	5,014	0.18	47.77	52.67	62.99	4,417,601
Price signal P2P bid offer (USD)	167,381	51.39	13.46	29.84	44.76	49.64	55.33	218.20
Price signal P2P ask offer (USD)	635,475	56.66	11.19	35.00	50.57	55.00	62.80	241.27
Price signal Buyback (USD)	777,038	51.14	13.99	0.00	49.12	50.79	54.42	185.80
AMM liquidity-intraday perspective								
AMM pool liquidity bid (USD)	776,532	1,144	2,913	0.00	47.36	847.14	1,577	89,069
AMM pool liquidity ask (USD)	776,532	1,144	2,913	0.00	47.36	847.14	1,577	89,069
Total AMM pool liquidity (USD)	776,532	2,287	5,826	0.00	94.71	1,694	3,155	178,139
AMM slippage (%)	555,169	-0.012	0.023	-1.000	-0.034	-0.005	-0.003	0.000
LP trading fee revenue (USD)	555,169	0.019	0.132	0.000	0.001	0.002	0.019	11.346
P2P liquidity-intraday perspective								
P2P bid offer depth (USD)	167,381	1,352	1,836	0.00	50.05	517.47	3,895	17,525
P2P ask offer depth (USD)	635,475	6,005	16,426	0.00	57.00	992.60	14,336	433,436
Total P2P offer depth (USD)	646,415	6,253	16,431	0.00	58.90	1,184	14,717	433,686
Time on market since offer creation (days)	207,885	17.89	48.40	0.00	0.00	1.81	48.16	1,140
Stress event variables-intraday perspective (± 100 -days around the shock)								
1-token price impact (bps)	235,404	1,640	2,501	-10,000	256.11	521.33	7,655	9,101
Bid-ask spread (bps)	64,912	1,472	730.97	0.00	605.60	1,393	2,438	4,919
Asset size (USD)	752	215,595	327,988	48,080	62,440	84,712	590,675	5,000,000
TRO affected	752	0.73	0.45	0.00	0.00	1.00	1.00	1.00
Market control variables at trade-daily perspective								
RWA token price volatility (7-day)	185,940	0.08	0.16	0.00	0.02	0.05	0.15	3.18
Market-to-appraisal ratio	188,193	1.06	0.05	0.93	0.99	1.06	1.11	1.19
Cumulative return ETH (one week before)	188,193	0.01	0.11	-0.52	-0.11	0.00	0.15	0.67
Average ETH tx fee (USD)	188,193	8.36	12.20	0.07	0.61	4.00	23.57	200.06
1m Treasury yield Δ	188,193	0.00	0.05	-0.55	-0.02	0.00	0.03	1.06
10y Treasury yield Δ	188,193	0.00	0.06	-0.30	-0.08	0.00	0.08	0.29
ADS Index Δ	188,193	-0.10	0.96	-26.69	-0.47	-0.12	0.36	9.58
S&P Case Shiller Index Δ	188,193	0.01	0.01	-0.01	-0.00	0.00	0.02	0.03

Note: This table reports the number of observations, mean, standard deviation, minimum, 10th percentile, median, 90th percentile, and maximum for the intradaily (transaction-level) and daily aggregated perspectives. All variables and the corresponding calculation are defined in Appendix Table A.5.

Trades and price dispersion. Roughly 40% of trades are purchases. Price signals across mechanisms differ substantially in dispersion. Buyback prices cluster tightly around appraisal values (mean = \$51.14). P2P bid-ask quotes bracket these appraisal values in a manner consistent with posted-offer trading. AMM-implied prices, by contrast, exhibit extremely high dispersion (SD = \$5,014) driven by near empty liquidity pools that mechanically generate outsized implied quotes when reserves are thin. Although median AMM prices align closely with P2P and Buyback prices, the mean is dominated by a small number of extreme outliers.

⁴The total token count is 3,150,484. RealT reports retaining 5% of tokens at primary issuance (157,524 retained tokens), leaving approximately 3 million tokens in free float.

AMM and P2P liquidity. Liquidity provision differs markedly across mechanisms. Observed pool liquidity per side ranges from essentially empty (min $\approx$ \$0) to deep pools (max = \$89,070), with central tendency around \$1,000 (median = \$847; mean = \$1,144, per side). AMM slippage is observed only when trades execute against the pool and is not imputed on AMM-inactive observations; conditional on AMM execution, slippage is modest on average (mean = -1.2%; median = -0.5%) but can become extreme when reserves are thin (min = -99.9%). LP fee revenues are correspondingly small for most trades (90th percentile = \$0.019; mean = \$0.019), reflecting the generally low dollar trade sizes settled on AMMs.

P2P liquidity is posted and therefore intermittent and not guaranteed. In the trade-matched sample, bid quotes are observed far less often than ask quotes (N = 167,381 vs. 635,475), consistent with greater standing sell interest. Median posted depth is \$517 on the bid side and \$993 on the ask side, but the distribution is highly right-skewed, with occasional very large ask-side inventories (max = \$433,436). The median time on market from offer creation until execution (partial or full fill) is roughly 2 days, indicating relatively rapid turnover; however, P2P execution remains probabilistic relative to the AMM’s immediate execution, and can take substantially longer (mean $\approx$ 18 days; 90th percentile $\approx$ 48 days).

Stress event variables. In the absence of a unique market-clearing reference price, we compare mechanisms using their implied execution prices for a fixed-size trade. *Price impact* is measured as the implied execution cost of a fixed-size trade and expressed in basis points. We evaluate price impact for a representative 1-token trade (\$50). For AMM trades, price impact is derived from the constant-product pricing rule using contemporaneous pool states. Let P^{spot} denote the AMM-implied spot price and L^{pool} the total liquidity pool value in USD. The implied base-token reserve is inferred as $X = L^{\text{pool}}/P^{\text{spot}}$. The price impact of executing a 1-token trade is computed as $10,000 \times (1 - \frac{X}{X+1})$, expressed in basis points, which captures the mechanical slippage induced by moving along the constant-product curve and reflects the execution cost implied by available AMM depth at the time of the transaction. For P2P trades, price impact is defined relative to the contemporaneous posted ask quote. Specifically, impact is computed as $10,000 \times (P^{\text{exec}} - P^{\text{ask}})/P^{\text{ask}}$, where P^{ask} is the liquidity-weighted ask price prevailing at the time of execution. This measure captures the price concession required to execute against posted bilateral liquidity.

In the later phase of the maturing platform starting with the end of 2024, we are able to directly observe two-sided bilateral quotes in the P2P market and construct a *bid–ask spread* as an alternative measure of market liquidity. For each P2P observation with both a realistic bid and ask quote, we compute the bid–ask spread as the midpoint-normalized difference between the prevailing liquidity-weighted ask and bid prices. Let P^{ask} and P^{bid} denote the contemporaneous P2P ask and bid quotes, respectively, and define the midpoint price as $P^{\text{mid}} = (P^{\text{ask}} + P^{\text{bid}})/2$. The bid–ask spread is then measured as $10,000 \times (P^{\text{ask}} - P^{\text{bid}})/P^{\text{mid}}$, expressed in basis points. This measure captures the instantaneous cost of immediacy in bilateral trading and reflects both quote competition and available depth on each side of the P2P market. Bid–ask spreads are computed only for observations with active two-sided quotes. *Temporary Restraining Order (TRO) affected* is an indicator equal to one if the token involved in the trade is subject to a temporary restraining order, and zero otherwise. *Asset size* measures the total concurrent USD appraisal value of the underlying real estate asset associated with the trade.

Market control variables. Lastly, we include a set of market controls capturing token-level

conditions, blockchain environment factors, and broader macroeconomic states that may affect execution and routing. These include short-horizon RWA price volatility, crypto-market conditions proxied by Ether returns and transaction fees, and macro-financial conditions measured by changes in U.S. Treasury yields, the ADS business conditions index, and lagged house price growth.

2.5 Empirical analysis

This section examines trader sorting across mechanisms and price dispersion (Section 2.5.1), liquidity provision (Sections 2.5.2), and trading frictions under systemic stress (Section 2.5.3).

2.5.1 Trader sorting across AMM, P2P, and Buyback

To test the sorting implications of Hypothesis 1, we estimate transaction-level sorting frameworks of the form

$$\Pr(\text{Trade}_{i,t} \in m) = \beta_1^{(m)} \log(\text{TradeVolume}_{i,t}) + \mathbf{C}'_{i,t} \boldsymbol{\gamma}^{(m)} + \mathbf{L}'_{i,t} \boldsymbol{\delta}^{(m)} + \alpha_i + \delta_t + \varepsilon_{i,t}^{(m)}. \quad (2.14)$$

where $\Pr(\text{Trade}_{i,t} \in m)$ is an indicator equal to one if transaction i on day t executes via mechanism $m \in \{\text{AMM}, \text{P2P}, \text{Buyback}\}$. The key regressor $\log(\text{TradeVolume}_{i,t})$ captures the size of the executed order. The vector $\mathbf{C}_{i,t}$ contains transaction-day execution cost controls, including blockchain fees, AMM slippage, short-term volatility, valuation deviations, ETH returns, and Treasury yield changes. The vector $\mathbf{L}_{i,t}$ captures mechanism-specific liquidity: AMM pool depth, P2P ask-side inventory, and the time-on-market of the active P2P listing. All regressions include token fixed effects (α_i) and year-month fixed effects (δ_t), and standard errors are clustered at the token level.

Table 2.3 reports six models. Columns 1–3 show baseline specifications that include only trade size and execution-cost controls. Columns 4–6 control these specifications with mechanism-specific liquidity.

Trade size is the dominant determinant of sorting. Across all specifications, trade size is the dominant determinant of mechanism choice. Larger trades are significantly less likely to execute via AMMs and more likely to route to P2P or Buyback. These sorting patterns directly reflect the mechanism-choice condition implied by Equation 2.3. Execution costs materially reallocate trades across mechanisms. Higher blockchain fees and higher slippage reduce AMM usage and increase routing toward P2P venues and Buyback, reflecting differences in fees, execution certainty, and price impact across mechanisms. Market-wide controls such as volatility, valuation deviations, and macro variables enter with intuitive signs but play a secondary role relative to trade size and execution frictions.

Mechanism-specific liquidity reinforces sorting patterns. Mechanism-specific liquidity significantly sharpens sorting patterns, greater AMM pool depth increases the probability of AMM execution and reduces routing to P2P and Buyback mechanisms. Conversely,

Table 2.3: Transaction-level mechanism choice

	AMM trade (1)	P2P trade (2)	Buyback trade (3)	AMM trade (4)	P2P trade (5)	Buyback trade (6)
log(Trade volume, USD)	-0.0631*** (0.0025)	0.0555*** (0.0025)	0.0076*** (0.0005)	-0.0606*** (0.0017)	0.0533*** (0.0016)	0.0090*** (0.0005)
Blockchain fee (USD)	-0.0079*** (0.0023)	0.0008 (0.0014)	0.0070*** (0.0016)	-0.0038*** (0.0005)	-0.0002 (0.0004)	0.0032*** (0.0005)
Market-to-appraisal ratio	0.1539*** (0.0554)	-0.1382** (0.0542)	-0.0158** (0.0077)	-0.1584*** (0.0607)	0.3180*** (0.0775)	0.0264* (0.0150)
Average ETH tx fee (USD)	0.0002*** (0.0001)	-0.0002*** (0.0001)	0.0000 (0.0000)	0.0006*** (0.0001)	-0.0007*** (0.0001)	0.0000 (0.0000)
AMM slippage (wavg)	-0.7786*** (0.0660)	0.6969*** (0.0607)	0.0817*** (0.0075)			
RWA token price volatility (7-day)	-0.0203 (0.0180)	0.0151 (0.0172)	0.0053*** (0.0020)	0.0095 (0.0139)	-0.0069 (0.0149)	0.0114*** (0.0043)
Cumulative return ETH (1 week)	-0.0610*** (0.0097)	0.0594*** (0.0097)	0.0016 (0.0013)	-0.0107 (0.0100)	-0.0051 (0.0117)	-0.0041 (0.0027)
1m Treasury yield Δ	-0.0034 (0.0260)	-0.0036 (0.0268)	0.0071* (0.0037)	-0.1095*** (0.0220)	0.0966*** (0.0233)	0.0120 (0.0076)
10y Treasury yield Δ	-0.0256** (0.0126)	0.0244* (0.0127)	0.0012 (0.0015)	-0.0443*** (0.0132)	0.0345** (0.0144)	0.0066* (0.0036)
ADS index Δ	0.0206 (0.0165)	-0.0126 (0.0158)	-0.0080** (0.0031)	-0.0096 (0.0213)	-0.0168 (0.1257)	0.0413 (0.0281)
AMM pool liquidity bid (kUSD)				0.0034*** (0.0006)	-0.0034*** (0.0010)	-0.0004** (0.0002)
P2P ask offer depth (kUSD)					0.0005*** (0.0001)	0.0000 (0.0000)
Time on market since offer creation (days)					0.0029*** (0.0001)	
Token FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	603,527	603,527	603,527	764,843	628,229	628,229
R ²	0.355	0.338	0.044	0.643	0.634	0.076
Within R ²	0.174	0.144	0.034	0.174	0.195	0.023

Note: This table reports transaction-level regressions of mechanism choice on trade characteristics, market conditions, and mechanism-specific liquidity measures. Cluster-robust standard errors (token level) are reported in parentheses. All regressions are estimated at the transaction level using linear probability models with token and year-month fixed effects. The main explanatory variable is the logarithm of the trade's USD notional value. Columns 1–3 report baseline specifications including log(Trade volume), blockchain fees, AMM slippage, market-to-appraisal ratios, ETH fees, rolling 7-day price volatility, ETH returns, Treasury yield changes, and ADS index changes. Columns 4–6 replace AMM slippage with mechanism-specific liquidity (scaled in thousand USD): AMM bid-side pool liquidity for AMM routing, P2P ask-side offer depth and time on market of the active listing for P2P routing, and both AMM bid-side pool liquidity and P2P ask offer depth. Note that AMM bid and ask side is symmetric given the constant-product AMM mechanism. *, **, *** denote significance at the 10%, 5%, and 1% levels.

deeper P2P ask-side inventory and longer time-on-market increase the likelihood of P2P execution, consistent with posted-offer trading and search frictions. Appendix A.5 shows that the observed routing patterns are not driven by the volume of quasi-arbitrage activity across mechanisms. Instead, trades sort across mechanisms in response to mechanism-specific execution costs and liquidity constraints that govern economically meaningful execution.

2.5.2 Liquidity provision across mechanisms

AMM liquidity provision. AMM pools provide continuously available two-sided quotes implied by the reserve ratio under the constant-product design. Liquidity is supplied through LP deposits into the pool’s token reserves and removed through withdrawals, which mechanically scale depth on both sides of the market. Table 2.4 (Panel A) summarizes AMM liquidity provision across pools, distinguishing issuer and external LP activity.

Across the full sample, we identify 275 unique external LPs, accounting for 4,597 deposits and 3,321 withdrawals. Issuer liquidity is concentrated at pool inception and remains largely stable thereafter: the number of issuer deposits ($N = 954$) closely matches the number of pools created and underlying RWA issuances, indicating limited active liquidity management over the pool life cycle.⁵

Figure 2.5 traces the combined issuer and external liquidity of all three AMMs over time. The March 2021 migration from Ethereum to Gnosis and the shift from Uniswap to Levinswap coincided with a marked increase in external LP entry. However, this initial expansion was short-lived, as Levinswap relied on a non-stablecoin quote token that collapsed in 2021, imposing substantial losses on external LPs and triggering a sharp contraction in external liquidity. Over the same period, the P2P mechanism gains volume share (Figure 2.3), reducing AMM trading volume and LPs’ expected fee revenue. External liquidity does not meaningfully recover thereafter, even after May 2022, when all newly created RWA pools were paired with USDC, reducing LPs’ inventory risk. External participation begins to rise again toward the end of 2023 and remains elevated through subsequent issuer regime changes. Overall, issuer-sponsored liquidity anchors AMM depth throughout the sample.

To empirically test these descriptive patterns of external liquidity provision and evaluate Hypothesis 2, we follow Lehar and Parlour (2025) and Zhu et al. (2025) and proxy AMM liquidity with USD pool size, often referred to as total value locked (TVL), but focus on the external LP component. We estimate the following pool-week regression:

$$y_{i,t+1} = \beta_1 x_{i,t}^{\text{fee}} + \beta_2 x_{i,t}^{\text{iss}} + \beta_3 v_{i,t} + \beta_4 x_{i,t}^{\text{iss}} \times d_{i,t} + \alpha_i + \delta_t + \varepsilon_{i,t}, \quad (2.15)$$

where i indexes pools and t indexes weeks. $y_{i,t+1}$ is the average external liquidity supplied to pool i in week $t+1$, measured in USD at prevailing pool prices. $x_{i,t}^{\text{fee}}$ denotes the total trading-fee revenue earned by pool i in week t , with fees fixed at 0.25% of volume and allocated pro rata across LPs. $x_{i,t}^{\text{iss}}$ is issuer-sponsored liquidity in pool i , valued at pool-implied prices


⁵A notable exception occurs in February 2024, when the issuer withdrew most of its per-pool contribution, reducing the seed position from 20 to 4 RWA tokens (approximately \$1,600 at mean token prices).

Table 2.4: Liquidity provision across market mechanisms

Statistic	N	Mean	St. Dev.	Min	10%	Median	90%	Max
<i>Panel A: AMM liquidity provision (event-level)</i>								
<i>Issuer LP</i>								
Deposit (USD)	954	1,181	1,536	1.56	101.48	207.76	2,075	18,939
Withdrawal (USD)	778	1,483	6,301	0.01	2.47	412.71	1,671	82,684
<i>External LPs</i>								
Deposit (USD)	4,597	225.50	1,233	0.00	3.08	85.11	456.98	49,383
Withdrawal (USD)	3,321	435.16	2,974	0.00	10.01	102.89	562.74	136,684
<i>Total liquidity activity</i>								
Deposit (USD)	5,551	389.67	1,339	0.00	4.58	101.12	1,275	49,383
Withdrawal (USD)	4,099	634.09	3,855	0.00	9.96	109.23	1,582	136,684
Protocol-internal activity (USD)	3,079	159.61	1,049	0.00	1.99	54.66	309.88	50,783
<i>Panel B: AMM liquidity provision (pool-week panel)</i>								
External liquidity (USD)	23,833	494.76	3,957	0.00	0.36	104.55	398.24	131,308
Issuer liquidity (USD)	23,833	1,580	3,292	0.00	40.02	1,259	2,423	73,341
LP trading fee revenue (USD)	23,833	0.38	2.23	0.00	0.01	0.08	0.49	141.50
AMM token exchange rate volatility	23,833	0.049	0.091	0.000	0.007	0.030	0.089	4.399
<i>Panel C: P2P two-sided quote provision (wallet-token pair level)</i>								
Bid and ask offer overlap (days)	835	23.17	39.60	0.00	1.18	9.01	67.61	651.96
Price signal bid offer (USD)	835	50.56	14.66	33.22	45.10	49.02	52.55	215.00
Price signal ask offer (USD)	835	56.11	17.22	40.74	50.18	52.65	61.48	228.99
Spread (simple, %)	835	11.04	8.28	-2.80	2.94	9.69	22.15	57.74
P2P bid offer depth (USD)	835	382.01	751.48	0.00	96.77	199.09	857.53	10,526
P2P ask offer depth (USD)	835	601.81	3,537	0.01	20.14	111.87	794.79	76,987
Total offer depth (USD)	835	185.54	403.54	0.00	14.26	100.10	300.30	6,217

Note: Panel A reports event-level AMM liquidity provision across all pools, split into *Issuer LP*, *External LPs*, and *Total*. Deposit and withdrawal amounts are measured as two-sided USD flows, computed as twice the USD change in the pool's quote-token reserve under the constant-product design. Activity initiated by the protocol's designated fee-collector (LevinMaker) wallet is reported separately. Panel B reports pool-week summary statistics computed on the weekly regression sample (restricted to pool-weeks with at least two swaps in week t and defined within-week volatility). Panel C summarizes *P2P two-sided quote provision* by wallet-token. Wallets qualify if realistic bid and ask offers overlap at least once and both sides have positive posted depth in at least one overlap window. Overlap is total bid-ask overlap duration (days). Bid/ask price signals are liquidity-weighted time-averaged posted prices (USD) over overlap windows; the spread is the percentage difference between the ask and bid price signals and can be negative if the bid exceeds the ask. Total offer depth is the smaller of the bid- and ask-side posted depth.

Figure 2.5: AMM liquidity: Issuer vs external providers



Note: This figure plots weekly total AMM depth decomposed into *issuer* and *external liquidity provision*. Since 2022, pools are paired exclusively with USDC (see Appendix A.2). The three major declines in AMM depth correspond to (i) the LEVIN/USD collapse in April 2021 (the then-dominant quote token), (ii) the ETH/USD drawdown during the Terra collapse in mid-2022 affecting (W)ETH-quoted pools, and (iii) RealT’s regime shift and near-complete liquidity withdrawal across all pools in February 2024.

and determining the subsidy to external LPs at the end of week t . $v_{i,t}$ captures pool-level price deviations associated with impermanent loss, measured as the standard deviation of log exchange-rate changes between the RWA token and the quote token, following (Capponi and Jia, 2021). The interaction term $x_{i,t}^{\text{iss}} \times d_{i,t}$ captures whether issuer-sponsored liquidity matters more for external liquidity provision when fee revenue is low (below its mean). We include pool fixed effects α_i and year-month fixed effects δ_t , which absorb time-invariant heterogeneity across the underlying RWAs and pool design choices (including the pool’s quote-token pairing) as well as common time-varying shocks (including platform-wide regime changes). Standard errors are clustered at the pool level. The summary statistics for the pool-week panel are presented in Table 2.4 (Panel B).

Table 2.5 presents our regression results. In the log-log specifications, coefficients can be interpreted as elasticities. For example, in the fixed-effects model (2), a 1% increase in trading-fee revenue at time t is associated with an approximately 0.66% increase in external liquidity at $t + 1$, holding other covariates and fixed effects constant. This indicates that higher trading-fee revenue predicts greater subsequent external liquidity, consistent with Hypothesis 2 and the view that fees constitute a primary component of LP profits. In contrast, higher token-price volatility is associated with lower external liquidity, consistent with greater potential losses and loss-versus-rebalancing risk (Milionis et al., 2024).

Issuer-sponsored liquidity is positively related to external liquidity. Economically, a larger issuer position increases overall depth, reduces expected slippage, and can attract more trading and fees. It can further serve as a commitment signal that increases perceived platform trustworthiness. Furthermore, it enlarges the subsidized income stream that external LPs receive through the issuer’s yield-generating RWA tokens in the pool. When trading fee revenue is low, as captured by the interaction term in the regression, external LPs rely more heavily on subsidies when endogenous trading fee income is insufficient. This pattern is consistent with the AMM liquidity provision framework’s comparative statics (Appendix A.2.1): since the

Table 2.5: External liquidity provision in AMM pools

	log(External liquidity _{t+1} , USD)	
	(1)	(2)
log(LP fee revenue _t , USD)	1.097*** (0.165)	0.656*** (0.119)
log(Issuer liquidity _t , USD)	0.683*** (0.034)	0.873*** (0.031)
log(Issuer liquidity _t , USD) × Low-trading-fee env.	0.094* (0.053)	0.166*** (0.037)
AMM token exchange rate volatility _t	-2.382*** (0.266)	-1.469*** (0.213)
Pool FE	No	Yes
Year-Month FE	No	Yes
Observations	23,833	23,833
R ²	0.299	0.371

Notes: This table reports pool-week regressions of $\log(\text{External liquidity}_{t+1})$, computed from end-of-week external liquidity provision in the subsequent calendar week. Trading fee revenue is the weekly sum of LP fees, and issuer liquidity is end-of-week issuer-provided liquidity. Volatility is the standard deviation of the post-swap log spot exchange rate (quote per base) across swap events within the pool-week. The low-trading-fee environment is proxied by the negative, mean-centered log trading-fee revenue (higher values = lower-fee weeks). Log denotes the natural logarithm; coefficients on logged regressors are interpretable as elasticities. Column (1) is pooled OLS and reports overall R^2 ; column (2) includes pool and year-month fixed effects and reports within- R^2 . Standard errors are clustered by pool. *, **, *** denote significance at the 10%, 5%, and 1% levels. Sample is restricted to pool-weeks with at least two swaps. Results are robust to alternative minimum swap-count thresholds per pool-week used to compute volatility and to estimation on a daily panel.


derivative of L' with respect to L_R decreases with γ , increases in L_R induce larger liquidity responses in low-trading fee revenue regimes.

The overall result, especially the positive coefficient of the LP fee revenues, points to a feedback loop: limited trading constrains fee revenue, which in turn limits external liquidity provision, preventing the endogenous deepening typically observed in crypto-native AMMs. Consistent with this self-reinforcing dynamic, AMM trading activity declines over time (Figures 2.3 and 2.4), with trading migrating toward P2P execution.

P2P liquidity provision. P2P liquidity is sparse and episodic, consistent with a bilateral posted-offer market. While ask-side depth is persistently available, bid-side depth increases only temporarily, most notably in early 2025, and then recedes. P2P liquidity provision, which we define as concurrent bid and ask quotes posted by the same wallet address, is even rarer. We identify only 55 “two-sided quote providers.” Table 2.4 (Panel C) summarizes their activity and Figure 2.6 places it in the context of the full P2P offer book. Two-sided quoting is generally short-lived (median overlap of 9 days) and shallow (median two-sided depth of \$100), consistent with opportunistic, non-inventory-based liquidity provision rather than continuous market making. Of these two-sided quote providers, 12 also provide liquidity in AMMs, and the issuer is never among them. Taken together, these patterns suggest little strategic external liquidity provision in P2P and limited overlap in liquidity provision across mechanisms.

Buyback. Buyback liquidity is deterministic and issuer-controlled. Weekly redemption caps per investor restrict total executable quantity, and Buyback does not generate continuous market-based price signals. We treat it as a separate liquidity mechanism at a stable appraisal-based valuation.

Figure 2.6: P2P liquidity: Two-sided quote providers



Note: This figure plots P2P liquidity measured as the total posted notional depth on the relevant side of the offer book (ask-side offers to sell and bid-side offers to buy), summed across all active listings on each day; unposted or hidden liquidity is not included. The lighter gray area isolates the portion of depth posted by wallets that simultaneously maintain bid and ask offers for the same token (*two-sided quote providers*)


Integrated order-book reconstruction. We combine AMM depth, P2P offer-book depth, and Buyback capacity into a unified order-book representation. Figure 2.7 shows the aggregated depth across mechanisms. AMM liquidity appears as a smooth but issuer-driven layer; P2P liquidity is sparse and asymmetric; and Buyback provides a stable backstop that activates intermittently.

Overall, RealT’s AMMs function less like a decentralized venue sustained by competitive liquidity provision and more like an issuer-maintained inventory facility. External LP participation exists but remains shallow relative to issuer-sponsored depth. While standard channels emphasized in crypto AMM studies, such as adverse selection, impermanent loss, and fee competition, also operate in this setting (Aquilina et al., 2024; Lehar et al., 2024), a key distinction is the persistent presence of a central initiator attempting to sustain external liquidity beyond launch. In most DeFi environments, AMM liquidity is initially seeded and then scales endogenously with trading volume and fee revenue, often supported by protocol-token incentives at launch (Heimbach et al., 2021; Capponi et al., 2024; Milionis et al., 2024; Lehar and Parlour, 2025). In our setting, by contrast, issuer-sponsored depth crowds in external LPs only to a limited extent. Low turnover constrains fee revenue, and issuer-provided yield subsidies on deposited RWA tokens generate liquidity that does not scale with market activity. External LP participation is therefore tilted toward harvesting issuer-subsidized income rather than supplying depth for market making.

The constraining factor is scale: even if all trades in our data were theoretically routed through AMMs at the observed 0.25% fee, aggregate LP revenue would remain modest, approximately \$103 thousand in total or about \$45 per active token-year, which at a typical price of \$50 and for balanced buy and sell orders translates into $\gamma = 0.9$.⁶ Formula (2.13)

⁶See Equation 2.12. We imply the total USD trading volume of \$41.2 million (see Table 2.1). Active token-years are token-years with at least one trade (2,258). USD per active token-year are computed as total theoretical fee revenue divided by 2,258.

Figure 2.7: AMM liquidity pool depth, P2P offer-book depth, and Buyback capacity



Note: This figure plots aggregate daily liquidity available across the three secondary-market mechanisms. AMM liquidity is measured as the total two-sided notional depth implied by constant-product pool reserves under pool-implied pricing. For each pool, we proxy two-sided depth as $2 \times$ the USD value of the quote-token reserve (valued at contemporaneous quote-token USD prices) and aggregate across all active pools each day. The total is mirrored at the x-axis for illustration purpose of ask and bid side. P2P liquidity is measured as the total posted notional depth on the relevant side of the offer book (ask-side offers to sell and bid-side offers to buy), summed across all active listings on each day; unposted or hidden liquidity is not included. Buyback capacity is derived from the issuer's deterministic redemption facility and is measured as the aggregate daily executable notional implied by platform-imposed quantity caps and prevailing appraisal-based Buyback prices. All series are constructed at the daily frequency and represent potential executable liquidity rather than realized trading volume. Light gray shaded areas mark periods in which only two mechanisms operate concurrently; in particular, the Buyback was suspended for most (though not all) tokens following the TRO.

illustrates the consequences. For realistic and observable capital costs and RWA yields of $r = 0.09$ and $\mu = 0.09$, the maximum profitable liquidity provision L' is 12 (resp. 20) tokens for issuer-provided liquidity of 5 tokens ($L_R = 5$) (resp. 15). In other words, as liquidity providers tend to stay far below these values, the traded volume so far is insufficient to support persistent depth in the absence of additional incentives.

As emphasized by Park and Stinner (2025), competition for liquidity in DeFi is intense. In tokenized RWA markets, permissioned participation can further limit liquidity mobility and shrink the set of eligible LPs. As a result, effective designs for sustaining AMM depth in inherently low-turnover, friction-constrained environments remain an important open challenge.

2.5.3 Difference-in-differences analysis: fragmentation under stress

AMM vs. P2P. The collapse of the TerraUSD (UST) stablecoin on May 7, 2022 provides a sharp, externally imposed shock to decentralized markets in DeFi.⁷ Importantly, 2022 is the only period in our sample in which AMM and P2P mechanisms are comparable in trading volume, while none of the tokens in our data are directly paired with UST or Luna, ruling out mechanical price transmission. Instead, the shock propagated through system-wide execution frictions, that are orthogonal to real-estate fundamentals of our observed RWA tokens.

We measure execution costs using fixed-size price impact (1 token or \$50), defined as the implied execution cost of trading a representative quantity of tokens under each mechanism. Price impact captures liquidity conditions implied by each mechanism's pricing rule and available depth. By holding the trade size fixed, this measure isolates mechanism-driven liquidity costs rather than endogenous sorting in realized trades. To quantify the effect of the TerraUSD collapse, we estimate the following difference-in-differences (DiD) specification:

$$Y_{i,t,m} = \beta_1 \text{AMM}_m + \beta_2 (\text{AMM}_m \times \text{Post}_t) + \mathbf{C}'_{i,t} \boldsymbol{\gamma} + \alpha_i + \delta_t + \varepsilon_{i,t,m}, \quad (2.16)$$

where $Y_{i,t,m}$ denotes the 1-token price impact (in bps) for token i on day t under mechanism m . The indicator AMM_m equals one for AMM executions, and Post_t equals one for observations after May 7, 2022. Token fixed effects α_i absorb all time-invariant heterogeneity across assets, while day fixed effects δ_t absorb aggregate shocks common to all tokens. Standard errors are clustered at the token level. Control variables include rolling 7-day price volatility and volume-weighted average transaction fees to capture time-varying execution conditions. The coefficient of interest, β_2 , identifies the differential change in AMM execution costs relative to P2P trading following the Terra collapse for the same token.

Table 2.6 reports the DiD estimates for the 1-token price impact. Column (1) shows results for the full sample. Even prior to the shock, AMM execution is associated with substantially higher implied price impact than P2P trading, reflecting thinner liquidity. Following the Terra collapse, AMM execution costs increase sharply. The interaction coefficient $\text{AMM} \times \text{Post}$ indicates an additional increase of approximately 506 basis points in the implied cost of

⁷At the time, UST was the third-largest stablecoin with a market capitalization of roughly \$18 billion. The collapse triggered ecosystem-wide volatility, liquidity withdrawals, fee spikes, and pricing distortions, including spillovers to other major stablecoins (Uhlig, 2022; Lee et al., 2023).

Table 2.6: 1-token price impact around the Terra collapse (AMM vs. P2P)

	1-token price impact (bps)				
	All tokens	Pool size split (USD)		Asset size split (USD)	
		Small pools ($<$ median)	Large pools ($>$ median)	Small assets ($<$ median)	Large assets ($>$ median)
	(1)	(2)	(3)	(4)	(5)
AMM	1,045.3*** (112.11)	2,598.6*** (368.87)	571.0*** (23.34)	2,618.2*** (319.63)	756.1*** (66.59)
RWA token price volatility (7-day)	73.0 (254.38)	400.9** (182.58)	-20.7 (53.19)	696.7** (314.95)	100.5 (286.41)
Total fees (USD, wavg)	2.1 (2.99)	0.3 (2.28)	0.2 (0.31)	-9.7 (6.01)	3.3 (2.07)
AMM $\times$ Post	506.0*** (152.63)	2,284.7*** (356.30)	82.7** (39.09)	598.1** (295.16)	216.8** (104.91)
Token FE	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes
Observations	235,404	71,506	140,982	78,745	156,659
R^2	0.085	0.353	0.436	0.154	0.128

Notes. This table reports DiD estimates for the 1-token price impact (bps) comparing execution costs on AMMs relative to P2P markets around the Terra collapse on May 7, 2022. Column (1) reports results for the full sample. Columns (2) and (3) split the sample by the median pre-event AMM pool size (USD). Columns (4) and (5) split the sample by the median underlying asset size (USD). All specifications include token and day fixed effects. Standard errors are clustered at the token level. ***, **, and * denote significance at the 1%, 5%, and 10% levels. We report within- R^2 . The results remain robust when all fixed-size price impact measures are winsorized at the 1st and 99th percentiles to limit the influence of extreme observation and the corresponding estimates are available in the replication package.


executing a 1-token trade on AMMs relative to P2P markets. Columns (2) and (3) split the sample by the median pre-event AMM pool size. The post-shock increase in AMM price impact is concentrated in small pools. For pools below the median size, the implied execution cost rises by more than 2,280 basis points, whereas for large pools the increase is economically much smaller, though still statistically significant. This monotonic attenuation across pool size directly implicates liquidity depth as the key state variable governing AMM resilience under stress. Columns (4) and (5) report analogous splits by underlying total asset size (USD). The post-Terra collapse increase in AMM execution costs is substantially larger for small assets than for large assets, consistent with lower turnover and thinner effective liquidity for smaller and more heterogeneous RWAs.

Figure 2.8 presents event-study estimates for the 1-token price impact. When aggregating across all pools, average effects evolve more smoothly around the Terra collapse. In contrast, event studies restricted to small pools (below median) show a sharp and persistent upward shift in execution costs around the shock date. Event-study estimates show no evidence of differential pre-trends in execution costs across mechanisms, supporting the parallel-trends assumption.

P2P vs. Buyback. On July 22, 2025, a Detroit court issued a Temporary Restraining Order (TRO) requiring rental collections for a subset of underlying properties to be redirected into escrow accounts until certificates of compliance were obtained.⁸ As a direct consequence, the issuer suspended the Buyback for all affected tokens. This legal intervention provides a plausibly exogenous shock to the issuer's redemption facility, allowing us to examine how the withdrawal of this backstop facility affects liquidity conditions in P2P markets, which at that

⁸For more details on the lawsuit, see <https://detroitmi.gov/news/real-token-properties-are-subject-sweeping-lawsuit>.

Figure 2.8: Event study of 1-token price impact around the Terra collapse



Notes: This figure plots difference-in-differences event-study estimates of AMM price impact relative to P2P markets around the Terra collapse on May 7, 2022. The outcome variable is the implied price impact of executing a trade of one token (\$50), expressed in basis points. Coefficients are normalized to zero on day -1 and are estimated from regressions including token and day fixed effects, as well as controls for rolling 7-day price volatility and volume-weighted average transaction fees and 95% confidence intervals based on standard errors clustered at the token level. The top panel reports results for all pools, while the bottom panel restricts the sample to tokens associated with small AMM pools (below median pool size in USD).

time were the dominant market mechanism. We estimate DiD regressions comparing P2P bid–ask spreads for tokens affected by the TRO to unaffected tokens before and after the event, with the following equation:

$$\text{Spread}_{it} = \beta (\text{TRO affected}_i \times \text{Post}_t) + \gamma_1 \text{Volatility}_{it} + \gamma_2 \text{Fees}_{it} + \alpha_i + \delta_t + \varepsilon_{it}, \quad (2.17)$$

where Spread_{it} denotes the midpoint-normalized P2P bid–ask spread (in bps) for token i on day t . The indicator TRO affected_i equals one for tokens whose Buyback were suspended due to the TRO, and Post_t equals one for observations on or after July 22, 2025. Token fixed effects α_i absorb all time-invariant heterogeneity across assets, including location, property characteristics, and baseline liquidity, while day fixed effects δ_t capture aggregate shocks common to all tokens. Standard errors are clustered at the token level. Control variables include rolling 7-day price volatility and volume-weighted average transaction fees to account for time-varying market conditions.


Table 2.7 reports the results. While on average, affected tokens exhibit no significant increase in P2P bid–ask spreads following the TRO, the effect reveals when splitting the sample by asset size (USD). Spreads for large assets (above median in USD) widen substantially after their Buyback is suspended. In contrast, spreads for small assets (below median in USD) decrease, consistent with a reallocation of trading activity into P2P markets once their Buyback is suspended, increasing quote competition conditional on two-sided liquidity. For smaller asset sizes, the Buyback crowds out P2P trading. This effect is concentrated on the bid side: the share of tokens with standing P2P bids increases only among small assets, while it remains essentially unchanged for large assets. The size-dependent effects are robust to controlling for a July 2, 2025 public disclosure listing the properties lacking certificates of compliance that were later subject to the TRO and issuer Buyback suspension.⁹

Figure 2.9 presents event-study estimates of P2P bid–ask spreads around the TRO. The estimates show no evidence of differential pre-trends. Following the TRO, bid–ask spreads for large assets increase persistently, indicating a sustained deterioration in liquidity. Spreads for small assets remain stable or decline modestly. The aggregate response masks these opposing effects across asset size groups.

Overall, the evidence supports Hypothesis 3 and shows that fragmentation amplifies trading frictions through interacting mechanism-specific channels. Under systemic stress, AMMs continue to clear trades mechanically, but rising price impact reduces trading incentives and fee revenue, weakening external liquidity provision and accelerating migration toward information-sensitive outside options. P2P markets, while slower to execute, absorb this flow and exhibit greater price stability under stress. The issuer-provided Buyback interacts asymmetrically with these dynamics. It anchors prices and mitigates dispersion for large assets, but crowds out bilateral trading for small assets in normal times. When the Buyback is suspended, P2P liquidity deteriorates for large tokens as centralized support is withdrawn, while participation and quote competition increase for small tokens as bilateral trading becomes the marginal liquidity source. Liquidity outcomes under stress therefore depend on asset scale and the interaction between execution mechanisms.

⁹The results are available in the replication package. For the July 2 filing/announcement document, see <https://outliermedia.org/wp-content/uploads/2025/07/City-v.-Real-Token-Complaint-1-1.pdf>.

Figure 2.9: Event study of P2P bid–ask spreads around the Temporary Restraining Order



Notes: This figure plots difference-in-differences event-study estimates of P2P bid–ask spreads around the Temporary Restraining Order (TRO) issued on July 22, 2025. The outcome variable is the midpoint-normalized P2P bid–ask spread, expressed in basis points. Coefficients are normalized to zero on day -1 and are estimated from regressions including token and day fixed effects, as well as controls for rolling 7-day price volatility and volume-weighted average transaction fees. Shaded areas represent 95% confidence intervals based on standard errors clustered at the token level. The top panel reports results for all assets, the middle panel for large assets (above median in USD), and the bottom panel for small assets (below median in USD), where asset size is defined based on the median underlying asset value in USD.

Table 2.7: P2P Bid–ask spreads around the Temporary Restraining Order (P2P vs. Buyback)

	Bid–ask spread (bps)		
	All tokens	Asset size split (USD)	
		Small assets ($<$ median)	Large assets ($>$ median)
	(1)	(2)	(3)
RWA token price volatility (7-day)	-39.30 (60.60)	54.12 (54.22)	-78.63 (83.00)
Total fees (USD, wavg)	-10.48*** (2.103)	-10.24*** (1.629)	-10.48*** (2.935)
TRO affected $\times$ Post	191.4 (117.2)	-193.9*** (61.99)	271.1** (123.5)
Token FE	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Observations	64,694	19,203	45,491
R^2	0.013	0.009	0.023

Notes: This table reports DiD estimates of P2P bid–ask spreads around the Temporary Restraining Order (TRO) issued on July 22, 2025, which led to the issuer-initiated suspension of the Buyback mechanism for 546 RWA tokens out of 752. The dependent variable is the midpoint-normalized bid–ask spread expressed in basis points. Treated tokens (*TRO affected*) are those for which Buyback was suspended due to the TRO. Column (1) reports results for the full sample; columns (2) and (3) split the sample by the median underlying asset size (USD). All specifications include token and day fixed effects and control for rolling 7-day price volatility and volume-weighted average transaction fees. Standard errors are clustered at the token level. ***, **, and * denote significance at the 1%, 5%, and 10% levels. We report within- R^2 . The results remain robust when all fixed-size price impact measures are winsorized at the 1st and 99th percentiles to limit the influence of extreme observations and the corresponding estimates are available in the replication package.

2.6 Conclusion

This paper studies liquidity formation in a tokenized market for illiquid RWAs using a uniquely transparent dataset of 777,038 secondary-market transactions. This fully on-chain environment is the only setting in which three fundamentally different liquidity mechanisms coexist under a uniform token design. The central implication is structural: fragmentation across heterogeneous mechanisms does not merely redistribute order flow, but shapes how liquidity is financed and sustained. When turnover is low or market conditions deteriorate, deterministic pricing remains operational but becomes fragile as execution costs rise and fee-based liquidity provision contracts, shifting activity toward alternative mechanisms in a manner that depends on asset scale and issuer backstop liquidity. Fragmentation therefore places sharp limits on when deterministic market design can sustain depth and becomes an amplifier of trading frictions rather than a neutral allocation of liquidity.

While tokenization can facilitate trading, the question of liquidity creation remains unresolved and highlights trade-offs in the secondary market design of inherently illiquid but tokenized RWAs. When platforms rely on AMMs to provide instant, two-sided liquidity, viability becomes highly scale-dependent. For small, slow-moving, and heterogeneous RWAs, thin pools generate high price impact, making rule-based pricing fragile and liquidity provision unlucrative. In addition, whitelisting, and especially token-level whitelisting, fragments participation and limits trading intensity and liquidity provision, further constraining sustainable liquidity. Buyback acts as a capped, appraisal-based liquidity backstop that absorbs larger redemptions at the cost of market-based price formation. P2P trading dominates economic volume due to lower effective price impact, but resembles classical OTC search markets with uncertain execution and largely hidden buy-side liquidity. This leaves unresolved who bears

the cost of visible, committed liquidity in fragmented tokenized markets for traditionally illiquid RWAs.

In conclusion, tokenization alone does not ensure liquid secondary markets for illiquid RWAs. When heterogeneous execution mechanisms coexist, liquidity formation is constrained by market design and incentive compatibility, leaving liquidity fragmented, fragile, and conditional on issuer support rather than fully endogenous. P2P sustains liquidity but resembles a digitized, technologically facilitated search market that does not guarantee continuous liquidity. Without improved mechanism design or deliberate choices about which frictions to tolerate, tokenized markets risk reproducing rather than resolving the liquidity constraints they seek to eliminate.

Appendix

A.1 RWA market and liquidity mechanisms

This section positions RealT within the broader landscape of RWA tokenization platforms. RealT’s joint use of liquidity-pool-based AMMs, bilateral P2P, and issuer redemption-style Buyback mechanisms, combined with fully transparent on-chain record-keeping, provides a comprehensive setting to study liquidity formation under the fragmentation often encountered in DeFi. In particular, the RealT ecosystem constitutes a natural laboratory for tokenized illiquid RWAs, offering uniquely rich transaction-level data across coexisting secondary-market mechanisms under a uniform token design. Table A.1 summarizes selected tokenized RWA platforms by underlying asset class and their overarching secondary-market liquidity mechanisms, with platform-specific implementations. Among the platforms summarized, RealT uniquely enables observation of all three prominent mechanism categories over an extended horizon, with the full history of trades and liquidity-relevant events recorded on public blockchains since 2019. Many other platforms rely on one or two of these same mechanism categories, making insights from the RealT setting naturally transferable.

Table A.1: Overview of RWA tokenization platforms and their liquidity mechanisms

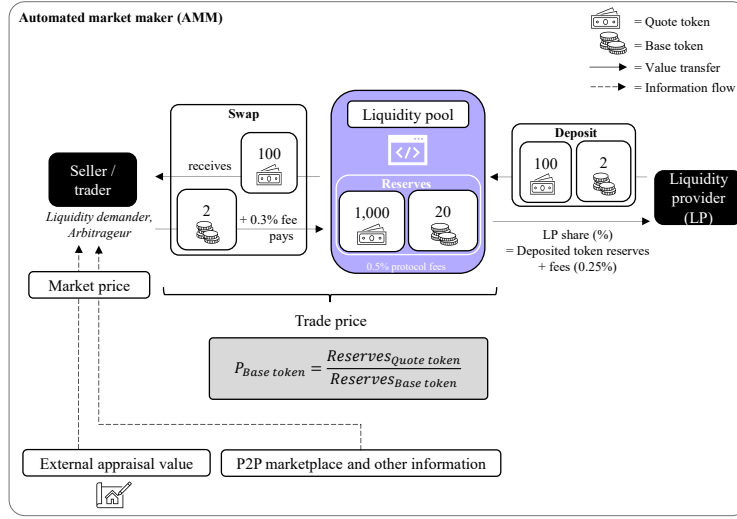
Platform	Asset class	Liquidity mechanism	Description
RealT	Real Estate	AMM/Pool-based (Uniswap, Levin-swap), P2P (issuer-non-operated), Buyback/Redemption (issuer)	Tokenized real estate platform using multiple DeFi-native liquidity mechanisms. High transparency, core case study for this paper.
Lofty.ai	Real Estate	P2P (centralized marketplace listings) / AMMs/Pool-based (Proactive Market Maker)	Real estate tokenization platform where individual issuers can tokenize properties and sell shares on a centralized secondary market, with valuation-based pricing and pool-based liquidity and P2P offer-book listings.
Reental	Real Estate	P2P (centralized marketplace listings; direct OTC transfers)	Real estate tokenization platform focusing on rental real estate and pre-developments in Spain and Latin America with platform-gated secondary trading.
Figure	Private Credit / HELOCs	AMM/Pool-based (yield/loan pools) / Alternative Trading Systems (ATS)	Publicly listed tokenization platform specialized in tokenized consumer loans, most notably Home Equity Lines of Credit (HELOCs). Provides ATS trading via Figure Markets.
Maple Finance	Private Credit	AMMs/Pool-based (underwritten lending pools, DeFi wrapper)	On-chain institutional private credit via curated lending pools. Liquidity is primarily via pool withdrawals; secondary exit liquidity is limited.
Centrifuge	Private Credit	AMMs/Pool-based (liquidity pools)	Tokenizes real-world credit, invoices and receivables for DeFi lending; decentralized credit pools.
Artory	Art	Secondary trading via Tokeny Integration	Partners with tokenization infrastructure providers to enable compliant secondary trading for traditionally illiquid art market on Ethereum.
Aktionariat	Private Equity / Stocks	On-chain order book / AMM/Pool-based (Hybrid)	Swiss platform allowing direct trading of tokenized company shares (AGs) using blockchain infrastructure.
Swarm Markets	Stocks / US Treasuries / Gold	AMM/Pool-based (Balancer-fork) / P2P ("dOTC", decentralized OTC)	German (BaFin)-regulated DeFi platform enabling compliant trading via AMM-based DEX and on-chain OTC service.
Securitize	Private Funds / Stocks	Centralized matching / Secondary market (ATS) / AMM/Pool-based (UniswapX)	SEC-registered platform offering third-party tokenized securities and funds with secondary trading via regulated ATS and DeFi integrations.
Ondo Finance	Stocks / ETFs / Bonds	Buyback/Redemption (24/5 redemption linked to public markets) / AMM/Pool-based	Institutional-grade RWA issuer offering tokenized Treasuries and listed stocks/ETFs with redemption options linked to underlying public markets and DeFi composability.

Note: This table provides a non-exhaustive overview of selected RWA tokenization platforms and the secondary-market liquidity mechanisms they employ. Platforms are classified into three broad mechanism categories: **AMM/Pool-based** mechanisms, which include AMMs with algorithmic pricing and continuous quoting and other pool-based capital allocation vehicles; **P2P** mechanisms based on bilateral or posted-offer trading; and **Buyback/Redemption** mechanisms offering issuer-provided liquidity at appraisal- or market-linked prices. Several platforms combine multiple mechanisms. Platform descriptions are based on publicly available documentation and observed market functionality.

A.1.1 RWA marketplaces and institutional details

AMM mechanism. AMMs (Figure A.1) provide continuous, two-sided liquidity through algorithmic pricing based on liquidity pools. Trades execute immediately at prices determined by the CPMM formula, implying convex price impact that increases with trade size. In RealT’s AMMs, LPs earn trading fees and continue to receive rental income on the RWA tokens they contribute. In addition, the issuer seeds each pool with retained tokens and redistributes the associated cash flows to external LPs pro rata. This creates a declining marginal subsidy: external LP yields are highest when the pool is small and fall as participation deepens.

Figure A.1: AMM mechanism



Note: This figure illustrates the schematic workflow of an automated market maker (AMM) mechanism.

Levenswap, the dominant AMM venue in our setting, implements the Uniswap V2 constant-product AMM design (Adams et al., 2020). AMMs offer high execution certainty and low fees but expose traders and LPs to slippage and inventory risk when liquidity is thin and paired pool tokens are uncorrelated. Arbitrageurs continuously adjust pool prices in response to order flow, linking AMM prices to other market signals (Angeris et al., 2019). Liquidity pools are typically created shortly after a RWA STO. Liquidity provision and withdrawal events in constant-product AMMs do not mechanically alter price signals, as reserves are adjusted proportionally on both sides of the pool.

AMM price signals and liquidity pools. We extract liquidity-pool reserve states before and after each executed token amount (swap) from Uniswap (V1/V2 on Ethereum) and Levenswap (Gnosis). Each RWA token is associated with a distinct liquidity pool, typically paired with a stablecoin quote token, most frequently USDC. Table A.2 reports the distribution of all quote-token regimes across the sample. The regimes are chronologically ordered from top to bottom: ETH and WETH on Ethereum until March 2021, LEVIN from March 2021 until its collapse later that year, WETH on Gnosis as an interim quote token, and USDC as the dominant quote token from May 2022 onward.

Pool-implied prices are computed using the CCPM rule, where the price equals the ratio of quote-token reserves to base-token reserves, i.e., $P = y/x$, with x denoting the reserve quantity of the RWA token (base-token) and y the reserve quantity of the quote token. This


Table A.2: AMM event counts by pool quote token

Pool quote token	N	Swap	Deposit	Withdrawal	Total
ETH	9	5,622	286	182	6,090
LEVIN	98	161,148	3,016	1,444	165,608
WETH	98	294,704	1,253	547	296,504
USDC	560	93,700	3,985	2,016	99,701
Total	765	555,174	8,540	4,189	567,903

Note: This table reports AMM event counts (*Swap*, *Deposit*, and *Withdrawal*) by pool quote token. *ETH* (Ether) is Ethereum’s native token, *WETH* (Wrapped Ether) is the ERC-20 wrapped representation of ETH used across Ethereum-compatible blockchains and applications, *LEVIN* is the Levinswap protocol token, and *USDC* is the USD-pegged stablecoin (USD Coin). *N* is the number of distinct liquidity pools using the respective quote token. The slight differing number of pools (765) and issued RWA tokens (752) is driven by partially duplicated pools from the Ethereum-to-Gnosis migration and by RWA tokens with pools but limited trade activity at the time of data collection.

construction yields a continuous AMM price signal that can be matched to AMM, P2P, and Buyback transactions. For AMM trades, we use the pre-swap reserve state implied by the executed token amounts to avoid mechanical endogeneity. Slippage is defined as the relative deviation between the execution price and the pre-swap pool price, measured as $(P_{\text{Execution}} - P_{\text{Before}})/P_{\text{Before}}$, where $P_{\text{Execution}}$ is the realized swap price and P_{Before} is the pool-implied price based on reserves immediately prior to execution.

Figure A.2: AMM: Levinswap web interface

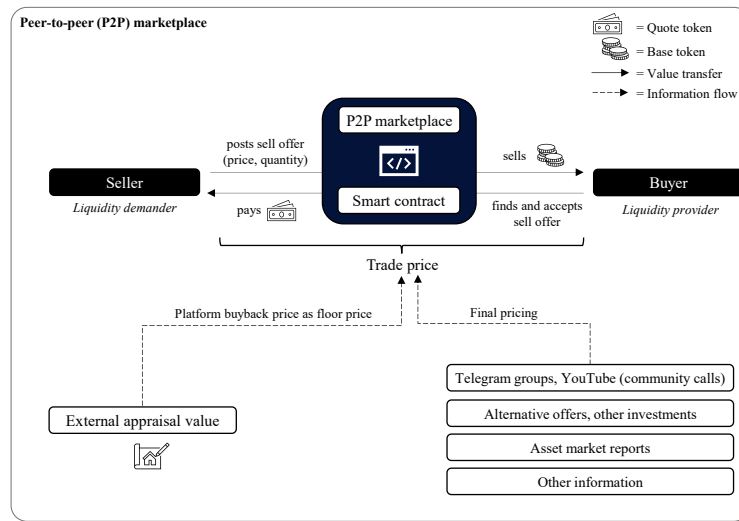


Note: The left and the middle screenshots show the Levinswap web interface for an AMM swap execution (trade), and the right screenshot shows the interface for liquidity provision. Accessed via <https://app.levinswap.realtoken.community/> on Oct 29, 2025.

P2P mechanism. P2P trading (Figure A.3) occurs via posted-offer marketplaces that allow users to specify prices and quantities directly. In our setting, trading takes place on the Swapcat and YAM marketplaces and almost exclusively on Gnosis. The marketplaces support buy, sell, and exchange orders through smart contracts. Offers become visible shortly after confirmation and remain active until matched, updated, or canceled. Execution is stochastic and depends on search frictions and posted depth, with limited buy-side liquidity and bid–ask spreads that remain persistent when two-sided quotes are posted.

YAM was introduced in December 2022 and provides a searchable live offer board with detailed metadata, including trade direction, price, quantity, implied yield, and the contemporaneous Buyback valuation. Private offers can be arranged here. Swapcat offers a simpler interface

Figure A.3: Peer-to-peer (P2P) marketplace



Note: This figure illustrates the transaction process on a smart-contract-based P2P market design.

without private offers. Both platforms record and settle offers and transactions on-chain. Overall, the observed P2P marketplaces resemble decentralized OTC markets with flexible pricing and uncertain execution.

Figure A.4: P2P: YAM web interface

> YAM Explore Your Offers Historic Sort...

Filters

Search

☐ Show only whitelisted properties's offers

☐ Hide dust offers

Sell Buy Exchange

Table

Offer ID	Offer Token	Buyer Token	Official Yield	Offer Yield	Yield delta	Official price	Price/token	Price delta	Available quantity
117580	Factoring PS-B-501	RealT RMT V3 USD	12.00%	11.81%	-0.78%	50 \$USD	50.4 (\$50.39)	+0.79%	3.999801587301587302
117611	16767 Greenfield	USD/C on xDai	11.83%	10.01%	-15.38%	50.6 \$USD	59.8 (\$59.79)	+18.17%	5.3
117617	2 Holdings Patton	USD/C on xDai	10.01%	10.01%	0.00%	50.21 \$USD	50.19 (\$50.18)	-	0.327852769769683203
117618	1 Holdings Pierson	USD/C on xDai	10.06%	9.20%	-8.58%	50.69 \$USD	55.45 (\$55.44)	+9.38%	3.243205
117657	13628 Tacoma	USD/C on xDai	11.26%	10.76%	-4.50%	50.61 \$USD	53 (\$52.99)	+4.71%	2
117664	9201 Longacre	USD/C on xDai	10.30%	10.30%	0.00%	52.02 \$USD	52.02 (\$52.01)	-	2
117665	13445 Tacoma	USD/C on xDai	10.56%	10.08%	-4.54%	52.5 \$USD	55 (\$54.99)	+4.75%	1.9162410450367773
117728	14126 Stout	USD/C on xDai	9.11%	8.35%	-8.33%	50.45 \$USD	55.04 (\$55.03)	+9.09%	0.192431
117730	14126 Stout	Wrapped XDAI	9.11%	8.35%	-8.32%	50.45 \$USD	55.04 (\$55.03)	+9.07%	0.192431
117731	14126 Stout	Realtoken Ecosystem USD	9.11%	8.35%	-8.34%	50.45 \$USD	55.04 (\$55.04)	+9.10%	0.192431

1 2 3 4 5 ... 829

Note: This figure illustrates the web interface to explore current offers on YAM. Accessed via <https://yam.realtoken.network/> on Oct 29, 2025.

P2P price signals and offer-book depth. P2P price signals are inferred from the reconstruction of the complete on-chain offer histories of the two active P2P marketplaces, Swapcat and YAM, and merged together on the token-level and under the application of specific realism filters, explained below. For each marketplace, we recover the full life cycle of every offer, including creation, updates, partial and full executions, and cancellations, using marketplace-specific parsing routines. On YAM, both prices and quantities can be updated, whereas on Swapcat only price updates are permitted. For offer creations, updates, and cancellations usual blockchain gas fees apply. Price signals are split into bid and ask quotes

based on offer direction, and reflect the contemporaneous, publicly visible offer inventory. For multiple offers for the same token at the same time, we apply the median and a liquidity-weighted aggregation, weighted by offer depth. Table A.3 reports daily summary statistics for realistic P2P depth and price measures at both the single-token and aggregate-market level. For each executed P2P trade, we compute *time on market* as the elapsed time between offer creation and execution. While this metric could alternatively be measured from the most recent price or quantity update, we consistently use *time on market since offer creation (days)* as the primary time-to-execution measure throughout the execution-cost framework and empirical analysis.

Realism filters. To ensure that P2P depth and price signals reflect executable and economically meaningful liquidity, we apply a sequence of realism filters. We exclude *private offers*, which are only available on YAM (3,951 offers), as they are not publicly observable.

We further remove offer states with zero *remaining inventory* (3,991 offers), as well as stale liquidity.

An offer-book state is classified as *stale* if its effective time on market, defined as the minimum of time since creation and time since last update, exceeds 60 days. Across marketplaces and sides, the 90th percentile of time since creation is approximately 49 days and the 90th percentile of time since last update is approximately 17 days, placing states persisting beyond 60 days in the extreme upper tail and unlikely to execute. In offer-state space, these screens remove 10,836 private states, 241,616 zero-inventory states, and 59,160 stale states.

Finally, we filter unrealistically priced offers by benchmarking P2P quotes against the concurrent Buyback price. The 90th percentile of the absolute price deviation is 24.6% on the ask side and 25.8% on the bid side; offers with deviations exceeding 30% are classified as *out-of-the-money* and removed (16,259 offers corresponding to 45,942 offer states).

On July 22, 2025, RealT temporarily suspended its Buyback for affected property tokens due to pending litigation; the same realism filters continue to apply during this period.

Overall, these filters remove approximately 14% of originally observed P2P offer-book depth when measured in USD inventory. Across the full sample of 777,038 transactions, this procedure yields 635,475 realistic ask-side and 167,381 realistic bid-side P2P price signals. The first P2P transaction occurs on October 26, 2020 on Swapcat, but activity increases substantially following the launch of YAM in early 2023. As a result, the majority of concurrent P2P price signals for non-P2P transactions are observed from 2023 onward (510,110 ask-side and 105,099 bid-side observations), with bid-side depth remaining sparse prior to the end of 2024, as illustrated in Figure 2.7.


Table A.3: Realistic P2P price signals order-book depth measures

Statistic	N	Mean	St. Dev.	Min	10%	Median	90%	Max
<i>One RWA token</i>								
Bid offer depth (USD)	1,671	186.74	455.96	0.00	0.00	0.00	574.89	3,034
Ask offer depth (USD)	1,671	932.89	1,156	0.00	58.00	597.04	2,029	15,069
Number of bid offers	1,671	1.25	2.50	0.00	0.00	0.00	5.00	14.00
Number of ask offers	1,671	3.04	2.18	0.00	1.00	2.50	6.00	12.00
Price signal bid offer (median, USD)	1,041	51.31	4.14	40.00	49.02	51.81	53.80	142.86
Price signal ask offer (median, USD)	1,629	56.04	7.87	45.99	52.46	55.42	58.00	168.00
Price signal bid offer (liquidity-weighted, USD)	1,041	51.38	4.10	40.00	49.16	51.87	53.88	142.86
Price signal ask offer (liquidity-weighted, USD)	1,629	56.43	7.84	45.99	53.02	55.88	58.46	168.00
<i>All RWA tokens</i>								
Bid offer depth (USD)	1,671	105,135	231,234	0.00	0.00	8,089	415,888	1,550,938
Ask offer depth (USD)	1,671	473,837	540,130	0.00	1,704	285,429	1,130,408	3,319,229
Number of bid offers	1,671	465.49	1,037	0.00	0.00	43.50	1,932	7,016
Number of ask offers	1,671	728.19	1,139	0.00	2.00	410.50	1,500	8,034
<i>Total activity counts</i>								
Total offers created								179,771
Total offer updates								127,867
Total purchases								207,885

Note: This table summarizes reconstructed realistic P2P price signals and offer-book depth after applying the realism filters defined in the accompanying Appendix text. For each token-day, offer-book measures are computed as the median across all observed offer events (offer creations, updates, purchases, and removals) within that token and day. *One RWA token* rows report summary statistics across days of the daily cross-sectional median across tokens. *All RWA tokens* rows report summary statistics across days of daily sums across tokens (depth and offer counts). Depth is measured in USD. Price statistics are computed only on days with at least one realistic quote on the corresponding side, so N may be below the number of calendar days.

Buyback mechanism. The Buyback mechanism (Figure A.5) provides a centralized issuer-provided liquidity outlet in which investors sell tokens back to the issuer at an appraisal-based price. Pricing is stable but infrequently updated, and execution is subject to quantity caps and deterministic delay (up to 10 working days). At RealT, the Buyback is capped at \$2,000 per wallet per week, requires that at least 95% of tokens are sold to investors, and settles with a deterministic waiting period of up to ten business days. Transactions settle on-chain in USDC via Ethereum or Gnosis and incur blockchain gas fees and a 3% platform redemption fee.

Figure A.5: Buyback



Note: This figure illustrates the Buyback process as a centralized market design.

Buyback price signals. Buyback price signals are based on third-party property appraisal values and retrieved from the RealT API. Using the full history of Buyback prices and their associated revaluation dates, we assign a contemporaneous Buyback price signal to every non-Buyback transaction. Because the Buyback mechanism was the first form of secondary trading available for these tokens, Buyback price signals are available without gaps across the entire non-Buyback transaction sample, yielding 777,038 observations.

STO process Our dataset consists of real estate tokens issued by RealToken LLC (RealT), a platform specializing in the tokenization of rental residential real estate. Under the Howey framework, these tokens qualify as investment contracts and are therefore regulated as securities under U.S. law. RealT initially conducted STOs under Regulation D 506(c) for U.S.-accredited investors, but as of October 29, 2025, all active real estate tokens have been converted to Regulation S, and U.S. investors are no longer permitted to invest in RealT real estate tokens.¹⁰ For further details on the STO process, see Kreppmeier et al. (2023).

Each tokenized property is held by a dedicated special purpose vehicle (SPV), RealToken Series LLC, which holds legal title to the property and operates as a standalone legal entity. Ownership interests in these SPVs are represented by ERC-20 tokens issued on Ethereum and Gnosis. The underlying assets are located predominantly in the United States with an increasing number abroad, mainly Latin America. Property management is predominantly outsourced to local operators. Investors acquire tokens during the STO after completing

¹⁰See <https://faq.realt.co/en/article/what-is-realt-who-can-invest-how-do-i-invest-1yyc5h5/>.

payment and digitally signing the offering memorandum, after which tokens are transferred to investor wallets via smart contracts. For most tokens, ownership represents an equity claim on the corresponding SPV, and net rental income, after operating expenses, insurance, and taxes, is distributed weekly to token holders through smart contracts. The RealT ecosystem further includes a decentralized autonomous organization (DAO) through which token holders can vote on specified governance matters, such as approving de-tokenizations following purchase offers for the underlying properties.

A.2 Stylized execution cost framework

Table A.4: Notation summary for the execution cost framework

Symbol	Definition	Specification / Notes
Prices and quantities		
P	Reference price (USD) for normalizing mechanism costs	Costs expressed relative to P
$q > 0$	Trade size (tokens), buy or sell, fractional allowed	Units of the base token
Fees		
F_A	AMM trading fee (relative)	Applied to notional qP
F_P	P2P trading fee (relative)	Applied to notional qP
$F_B = \frac{\Delta P_B}{P}$	Buyback discount relative to reference price	$\Delta P_B = P - P_B$; captures appraisal-based haircut plus redemption fee charged by issuer
Liquidity variables		
L_A	AMM liquidity (base-token units in pool)	Determines AMM slippage via constant-product rule
L_P	P2P liquidity (token units on relevant book side)	Ask depth for buys, bid depth for sells
Time and preference parameters		
τ_B	Expected time to execution in Buyback	Deterministic waiting time
$c\tau_B$	Expected P2P execution time (time on market) when $L_P = 0$	$c \neq 1$ allows faster/slower matching than Buyback
$\lambda \in [0, 1]$	Time preference / urgency parameter	Higher λ penalizes delay more strongly
Price-impact terms		
PI^0	Expected P2P price impact when $L_P = 0$	Bounded above by F_B ; thin-market concession
PI^+	Expected P2P impact when $q = L_P > 0$	Benchmark for depth-based impact calibration (“walking the book”)
$\delta \in (0, 1)$	Curvature parameter for sub-quadratic P2P impact	Implies impact exponent $1 + \delta \in (1, 2)$

Note: This table summarizes the notation used in the mechanism choice. All costs are expressed in units of qP and therefore normalized with respect to the reference price P . P2P liquidity L_P refers to the standing depth on the relevant side of the offer book (ask-side offers to sell and bid-side offers to buy). The impact parameters PI^0 , PI^+ , and δ govern price concessions in thin or shallow P2P markets, while AMM slippage is determined by the constant-product pricing rule through L_A . Time-to-execution enters linearly in expected delay costs through λ . All functional forms below should be interpreted as reduced-form approximations calibrated directly from the data (fees, depth, execution time) rather than structural primitives.

Trader preferences and choice.

Following conventions and DeFi AMM studies (e.g., Capponi and Jia (2021)), we treat the RWA token as the base asset or token and the stablecoin as the quote asset or token. Prices are quoted in units of the quote asset.

Mechanism-specific execution costs. All three mechanisms share the same reference

price P but differ in fee schedules, price impact, and time-to-execution. Each total cost expression contains the factor $qP > 0$; when comparing mechanisms, we divide by qP and work with normalized costs.

AMM. For an AMM with base-token depth L_A and reference price P , the post-trade price after a trade of size q is

$$P_{\text{new}} = \begin{cases} P \left(\frac{L_A}{L_A - q} \right)^2 & (\text{buy}), \\ P \left(\frac{L_A}{L_A + q} \right)^2 & (\text{sell}). \end{cases} \quad (2.18)$$

The relative slippage is therefore $\left(\frac{L_A}{L_A - q} \right)^2 - 1$ for buys and $1 - \left(\frac{L_A}{L_A + q} \right)^2$ for sells. Since the trader executes along the AMM price curve, we proxy the average execution-price impact by one half of the marginal price change from P to P_{new} , which is the convention in AMM trade settlements but also holds approximately when settling a large trader through of various small orders in P2P markets. Normalized AMM costs are

$$\begin{aligned} c_A^{\text{buy}}(q; L_A) &:= \frac{C_A^{\text{buy}}}{qP} = F_A + \frac{1}{2} \left[\left(\frac{L_A}{L_A - q} \right)^2 - 1 \right], \\ c_A^{\text{sell}}(q; L_A) &:= \frac{C_A^{\text{sell}}}{qP} = F_A + \frac{1}{2} \left[1 - \left(\frac{L_A}{L_A + q} \right)^2 \right]. \end{aligned} \quad (2.19)$$

The fee term $F_A qP$ is an approximation, because the actual fee depends on the realized transaction price. The error is second-order relative to slippage and empirically negligible in our setting.

Buyback. The Buyback executes at a fixed discount F_B but with deterministic delay τ_B . We assume delay costs scale with invested volume qP , so the normalized Buyback cost is

$$c_B := \frac{C_B}{qP} = F_B + \lambda \tau_B. \quad (2.20)$$

P2P with no standing depth ($L_P = 0$). When the relevant side of the P2P book is empty, the trader faces search frictions and a thin-market concession. We summarize the expected price concession as PI^0 and the expected execution time as $c\tau_B$ (with $c \neq 1$ allowing for faster or slower matching than the Buyback). Normalized P2P costs are

$$c_P^{(0)} := \frac{C_P^{(0)}}{qP} = F_P + PI^0 + \lambda c\tau_B. \quad (2.21)$$

A natural upper bound on PI^0 in case of a sell order is the Buyback discount F_B .

P2P with depth ($L_P > 0$). When there is standing depth $L_P > 0$ on the relevant side of the P2P book, we assume sub-quadratic price impact, consistent with the empirical literature (Almgren and Chriss, 2001) and market-microstructure intuition of “walking the book”:

$$c_P^{(+)}(q; L_P) := \frac{C_P^{(+)}}{qP} = F_P + \frac{1}{2} PI^+ \left(\frac{q}{L_P} \right)^\delta, \delta \in (0, 1). \quad (2.22)$$

This implies impact $\propto q^{1+\delta}$, with exponent between linear and quadratic.

Routing for buy orders.

For buy orders, the Buyback does not exist in our empirical setting, so we focus on AMM vs. P2P. The trader compares c_A^{buy} and $c_P^{(+)}$ when $L_P > 0$ and c_A^{buy} and $c_P^{(0)}$ when $L_P = 0$.

Main case: $L_P > 0$.

Proposition 1. A buy order of size q is routed to P2P (with $L_P > 0$) if

$$\left(\frac{L_A}{L_A - q}\right)^2 - 1 - PI^+\left(\frac{q}{L_P}\right)^\delta > 2F_P - 2F_A, \quad (2.23)$$

and to the AMM otherwise.

Minor case: $L_P = 0$.

Proposition 2. A buy order of size q is routed to P2P (with $L_P = 0$) if

$$\lambda < \frac{1}{c\tau_B} \left(\frac{1}{2} \left[\left(\frac{L_A}{L_A - q} \right)^2 - 1 \right] + F_A - PI^0 - F_P \right), \quad (2.24)$$

and to the AMM otherwise.

Routing for sell orders.

For sell orders, all three mechanisms are potentially relevant. The trader compares c_A^{sell} , c_B , and $c_P^{(0)}$ or $c_P^{(+)}$ depending on L_P .

Condition 1. Given a sell request (q, λ) , the AMM is preferred to the Buyback if

$$\lambda < \frac{1}{\tau_B} \left(F_A - F_B + \frac{1}{2} - \frac{1}{2} \left(\frac{L_A}{L_A + q} \right)^2 \right). \quad (2.25)$$

Condition 2. Given a sell request (q, λ) and $L_P = 0$, P2P is preferred to Buyback if

$$\lambda < \frac{1}{\tau_B(c-1)} (F_B - PI^0 - F_P) \quad \text{if } c > 1, \quad (2.26)$$

and

$$\lambda > \frac{1}{\tau_B(c-1)} (F_B - PI^0 - F_P) \quad \text{if } c < 1. \quad (2.27)$$

Condition 3. Given a sell request (q, λ) and $L_P > 0$, P2P is preferred to Buyback if

$$\lambda > \frac{1}{\tau_B} \left(F_P - \frac{1}{2} PI^+\left(\frac{q}{L_P}\right)^\delta - F_B \right). \quad (2.28)$$

Condition 4. Given a sell request (q, λ) and $L_P = 0$, P2P is preferred to the AMM if

$$\lambda < \frac{1}{c\tau_B} \left(F_A + \frac{1}{2} - \frac{1}{2} \left(\frac{L_A}{L_A + q} \right)^2 \right). \quad (2.29)$$

Condition 5. Given a sell request (q, λ) and $L_P > 0$, P2P is preferred to the AMM if

$$1 - \left(\frac{L_A}{L_A + q} \right)^2 - PI^+ \left(\frac{q}{L_P} \right)^\delta > 2F_P - 2F_A. \quad (2.30)$$

The inequalities in Conditions 1–5 partition (q, λ) -space into regions where AMM, P2P, or Buyback is optimal, conditional on $(L_A, L_P, F_A, F_P, F_B, \tau_B, c, PI^0, PI^+, \delta)$.

Comparative statics in liquidity

Proposition 3. All else equal, AMM becomes more preferable over P2P when L_A increases, while P2P becomes more attractive when L_P increases.

Proof. Consider the buy case with $L_P > 0$ and define

$$\Delta(L_A, L_P; q) := \left(\frac{L_A}{L_A - q} \right)^2 - 1 - PI^+ \left(\frac{q}{L_P} \right)^\delta. \quad (2.31)$$

Then

$$\frac{\partial}{\partial L_A} \left(\frac{L_A}{L_A - q} \right)^2 = \frac{-2L_A q}{(L_A - q)^3} < 0 \quad (2.32)$$

for $L_A > q$, and


$$\frac{\partial}{\partial L_P} \left[-PI^+ \left(\frac{q}{L_P} \right)^\delta \right] = PI^+ \delta q^\delta L_P^{-\delta-1} > 0. \quad (2.33)$$

Thus $\partial\Delta/\partial L_A < 0$ and $\partial\Delta/\partial L_P > 0$. Since P2P is chosen whenever $\Delta(L_A, L_P; q) > 2F_P - 2F_A$, a c.p. increase in L_A or a c.p. decrease in L_P lowers the advantage of P2P and shifts routing toward the AMM. On the sell side, the corresponding term replaces $\left(\frac{L_A}{L_A - q} \right)^2 - 1$ with $1 - \left(\frac{L_A}{L_A + q} \right)^2$; the derivative signs are unchanged.


Illustrative example

Figures A.6 and A.7 illustrate the implied routing regions for realistic parameter values, which can be observed in our data: $F_B = 0.03$, $F_A = 0.003$, $F_P = 0$, $L_P = 3$ (if not 0), $L_A = 20$, $\tau_B = \frac{1}{36}$ (10 days), $c = 1.8$, $PI^0 = 0.025$, $PI^+ = 0.12$, and $\delta = 0.5$.

Figure A.6: Buy order: rational choice between AMM and P2P (left: $L_P = 0$, right: $L_P > 0$)



Note: Dashed lines indicate a decrease in L_A by two tokens.

Figure A.7: Sell order: rational mechanism choice (left: $L_P = 0$, right: $L_P > 0$)


Note: Dashed lines indicate a decrease in L_A by two tokens.

A.2.1 AMM liquidity provision framework

In each liquidity pool, the RWA token is the base asset and a cryptocurrency token (typically a USD stablecoin) is the quote asset. The token initiator seeds the pool with L_R base tokens and the corresponding amount of the quote token at price P (initially equal to the Buyback price). An external LP chooses how many additional tokens $L > 0$ to contribute, along with the matched quote amount, so that the total base-token depth in the pool is $L_R + L$.

The LP keeps this position in the AMM for a fixed horizon (e.g. one year). Over this horizon, the pool faces a total sell volume q_s (base tokens sold into the pool) and buy volume q_b (base tokens bought out of the pool). We define

$$\Delta q := q_b - q_s, \quad Q := q_b + q_s. \quad (2.34)$$

Here, Δq is the net order imbalance (positive when buy pressure dominates), and Q is the total trading activity. At the end of the period, the base-token inventory in the pool equals $L_R + L - \Delta q$. Ex ante, Q and Δq are unknown and must be forecast by the LP.

The LP's profit over the horizon has four components:

1. Rental income on the tokens in the pool, i.e. the yield from the token held in the pool ($L\mu P$) plus the sponsored yield from the token initiator ($L_R\mu P$), resulting in the total yield component of:

$$\left(L + L_R - \frac{\Delta q L}{2(L_R + L)} \right) \mu P \quad (2.35)$$

(the subtrahend corrects for the varying number of tokens held during the interval, μ denotes the rental yield).

2. Trading fee income proportional to the total traded number Q amounting to

$$Q F_A \frac{L}{L_R + L} P. \quad (2.36)$$

3. The valuation changes arising from the constant-product pricing rule. Initially an amount of capital of $2P(L_R + L)$ is invested in the liquidity pool, which after the period amounts to

$$2P(L_R + L) \frac{L_R + L}{L_R + L - \Delta q}. \quad (2.37)$$

Therefore, for the LP the profit/loss from this valuation change is:

$$2P(L_R + L) \frac{\Delta q}{L_R + L - \Delta q} \frac{L}{L + L_R} = 2PL \frac{\Delta q}{L_R + L - \Delta q}. \quad (2.38)$$

This corresponds to the LP's fraction of the total profit in the liquidity pool. However, for reasons of mathematical tractability, we first approximate this term by linearizing it

$$2PL \frac{\Delta q}{L_R + L}. \quad (2.39)$$

Note that for small values of Δq (against $L_R + L$), which are realistic in our context, the approximation makes only a negligible error.

4. The opportunity cost of capital used to finance the AMM position.

$$-2rPL. \quad (2.40)$$

Collecting terms and normalizing by P , the expected profit can be written as

$$\Pi(L) = \frac{aL^2 + bL + c}{L_R + L}, \quad (2.41)$$

with

$$\begin{aligned} a &= \mu - 2r, \\ b &= 2L_R(\mu - r) + QF_A + \Delta q \left(2 - \frac{\mu}{2}\right), \\ c &= \mu L_R^2. \end{aligned} \quad (2.42)$$

where μ is the yield earned by holding the base token, r is the LP's cost of capital, and F_A is the AMM fee rate. We focus on the empirically relevant case $2r > \mu$, under which $\Pi(L)$ is concave in L .

External AMM liquidity

Let the “revenue intensity” be $\gamma := QF_A + \Delta q \left(2 - \frac{\mu}{2}\right)$. If $2r > \mu$ and $\gamma > 0$, the LP's profit-maximizing external liquidity is

$$L^* = -L_R + \sqrt{L_R \frac{\gamma}{2r - \mu}}. \quad (2.43)$$

Obviously, optimal AMM depth increases with γ , which captures expected fee income QF_A and expected gains from net buy imbalance Δq . When expected revenues are low relative to the capital cost spread $2r - \mu$, optimal external liquidity is small or in many realistic cases even negative, i.e. the optimum value cannot be achieved and the marginal optimum is

$L = 0$. Thus, for empirical implementation, it is more relevant to characterize the range of L for which liquidity provision is profitable at all.

Profitability region and the role of L_R

Suppose $2r > \mu$ and define γ as above. Then analyzing the numerator of (2.41) with respect to zeros is a straightforward quadratic problem resulting in the LP's profit being non-negative for all $L \in [0, L']$ with

$$L' = \frac{1}{2(2r - \mu)} \left[2L_R(\mu - r) + \gamma + \sqrt{4L_R^2 r^2 + 4L_R(\mu - r)\gamma + \gamma^2} \right]. \quad (2.44)$$

Moreover, $\partial L' / \partial L_R > 0$ implies higher initiator seed liquidity L_R strictly enlarges the range of profitable external liquidity provision. Note that for $r = \mu$ and an estimated $\gamma = 0$ (which can approximate a thin trading activity and balanced buy and sell orders), we have $L' = L_R$.

Equation (2.44) delivers a direct link between initiator commitment and sustainable AMM depth: the more tokens the initiator seeds, the larger the set of $(Q, \Delta q)$ realizations for which external LPs can earn non-negative profits. In our empirical tests, this prediction translates into higher observed AMM depth and higher AMM usage when L_R is high.

Differentiating (2.44) with respect to L_R yields

$$\begin{aligned} \frac{\partial L'}{\partial L_R} = \frac{1}{2(2r - \mu)} & \left[2(\mu - r) \right. \\ & \left. + \frac{4L_R r^2 + 2(\mu - r)\gamma}{\sqrt{4L_R^2 r^2 + 4L_R(\mu - r)\gamma + \gamma^2}} \right]. \end{aligned} \quad (2.45)$$

If we define $D := 4L_R^2 r^2 + 4L_R(\mu - r)\gamma + \gamma^2$, then

$$\frac{\partial L'}{\partial L_R} = \frac{2(\mu - r)\sqrt{D} + 4L_R r^2 + 2(\mu - r)\gamma}{2(2r - \mu)\sqrt{D}}. \quad (2.46)$$

Here, the denominator is positive because of $2r > \mu$. To see the positivity of the numerator of (2.46), consider that D can be written as

$$D = (\gamma + 2L_R(\mu - r))^2 + 4L_R^2 \mu(2r - \mu).$$

Thus $\sqrt{D} \geq |\gamma + 2L_R(\mu - r)|$, which implies (as a weaker inequality) $\sqrt{D} + \gamma \geq -2L_R(\mu - r)$. The numerator of (2.46) can be rewritten as

$$N := 2(\mu - r)(\sqrt{D} + \gamma) + 4L_R r^2.$$

Using the inequality regarding $\sqrt{D} + \gamma$, we have

$$N \geq 2(\mu - r)(-2L_R(\mu - r)) + 4L_R r^2 = 4L_R \mu(2r - \mu) > 0 \quad (2.47)$$

Altogether the claim $\partial L' / \partial L_R > 0$ is proven.

A.3 Definition of all variables

Table A.5: Definition of all variables

Variable	Definition	Calculation
Secondary-market variables (intraday perspective)		
Trades and price signals		
<i>Trade size (RWA token)</i>	Absolute number of RWA tokens transferred in the trade (base-token amount; for AMM trades: the exact token amount swapped in the pool).	Own
<i>Trade volume (USD)</i>	USD notional of the trade (quote-leg value): the quote-token amount exchanged, converted to USD at the trade timestamp (for AMM trades: the exact quote-token amount swapped in the pool, valued in USD).	Own
<i>Total fees (USD)</i>	Total fees paid in a trade, converted to USD: blockchain transaction fees (all mechanisms), issuer-platform fee on Buyback trades (3% of volume), and AMM trading fees on AMM trades (0.3% of volume, split into 0.25% to LPs and 0.05% to the protocol).	Own
<i>Trade price (USD/RWA token)</i>	Execution price per RWA token, expressed in USD. For AMM trades, computed from on-chain swap amounts and reserve states.	Own
<i>RWA token buy</i>	Dummy equal to one for buyer-initiated trades (purchase of RWA token), and zero for seller-initiated trades.	Own
<i>Ethereum blockchain</i>	Dummy equal to one if the trade executes on Ethereum, and zero if on the Gnosis.	Own
<i>Price signal Buyback (USD)</i>	Buyback price (USD per RWA token) quoted by the issuer platform at the time of the trade.	RealT
<i>Price signal AMM (USD)</i>	AMM-implied price (USD per RWA token) computed from pool reserves at the time of a P2P/Buyback trade, or immediately prior to an AMM trade.	Own
<i>Price signal P2P bid offer (USD)</i>	Posted bid price (USD per RWA token), i.e., willingness to buy the RWA token for a posted quote-token amount. If multiple bids are available at trade time, the median bid price is used.	Own
<i>Price signal P2P ask offer (USD)</i>	Posted ask price (USD per RWA token), i.e., willingness to sell the RWA token for a posted quote-token amount. If multiple asks are available at trade time, the median ask price is used.	Own
AMM and P2P liquidity		
<i>AMM pool liquidity bid (USD)</i>	USD value of the quote-token reserve in the AMM pool at trade time (capacity to buy RWA tokens).	Own
<i>AMM pool liquidity ask (USD)</i>	USD value of the RWA-token reserve in the AMM pool at trade time (capacity to sell RWA tokens), valued at the pool-implied price.	Own
<i>Total AMM pool liquidity (USD)</i>	Total AMM pool value (USD) at trade time, defined as the sum of the USD values of both reserves.	Own
<i>Deposit (USD)</i>	Quote-token amount deposited into the AMM liquidity pool (USD), paired with an RWA-token deposit in the prevailing reserve ratio.	Own
<i>Deposit (RWA token)</i>	RWA-token amount deposited into the AMM liquidity pool, paired with a quote-token deposit in the prevailing reserve ratio.	Own
<i>Withdrawal (USD)</i>	Quote-token amount withdrawn from the AMM liquidity pool (USD) as part of a paired liquidity removal.	Own
<i>Withdrawal (RWA token)</i>	RWA-token amount withdrawn from the AMM liquidity pool as part of a paired liquidity removal.	Own
<i>Swap (USD)</i>	Quote-token amount swapped in the AMM (USD, input amount), conditional on a trade executing against the pool.	Own
<i>Swap (RWA token)</i>	RWA-token amount swapped in the AMM (input amount), conditional on a trade executing against the pool.	Own
<i>Issuer liquidity (USD)</i>	USD value of the AMM pool liquidity attributable to the issuer-designated LP address, identified as the pool-creation (first-seeder) wallet across all pools.	Own
<i>LP trading fee revenue (USD)</i>	Liquidity-provider fee paid on AMM trades (0.25% of trade volume), allocated pro rata to LPs by liquidity share (with an additional 0.05% protocol fee to protocol-token holders).	Own
<i>AMM token exchange rate volatility</i>	Standard deviation of log changes in the AMM-implied exchange rate between the RWA token and the quote token.	Own
<i>AMM slippage (%)</i>	Percentage difference between the pool-implied pre-trade price and the realized execution price on the AMM; computed from on-chain swap amounts and reserve states and therefore reflecting trading fees.	Own
<i>P2P bid offer depth (USD)</i>	Aggregate posted bid depth, converted to USD for a given RWA token at the time of the trade (standing willingness to buy).	Own
<i>P2P ask offer depth (USD)</i>	Aggregate posted ask depth, converted to USD for a given RWA token at the time of the trade (standing willingness to sell).	Own
<i>Total P2P offer depth (USD)</i>	Total posted depth, converted to USD for a given RWA token at the time of the trade, summed across bid and ask sides.	Own
<i>Time on market since offer creation (days)</i>	Elapsed time (days) between offer creation and execution (partial or full fill).	Own
<i>Bid and ask offer overlap (days)</i>	Number of days in which the same wallet simultaneously posts at least one bid and at least one ask offer for the same RWA token.	Own
<i>Spread (simple, %)</i>	Simple percentage spread based on the overlapping ask and bid prices, computed as the percentage difference between the ask price and the bid price. The measure can be negative if the selected bid exceeds the selected ask.	Own

<i>Total offer depth (USD)</i>	Two-sided liquidity proxy for a market-maker–token pair, defined as the smaller of average bid offer depth and average ask offer depth during bid–ask overlap periods.	Own
Quasi-arbitrage opportunities and realized quasi-arbitrage trades		
<i>AMM → P2P bid</i>	Quasi-arbitrage opportunity implied by cross-mechanism price signals and defined as an economically feasible ordered pair of buying a RWA token via the AMM and selling into the P2P bid. Importantly, these opportunities are inferred from price signals and need not be feasible at economically meaningful scale or implementable for all investors due to execution risk, quantity caps, and whitelisting.	Own
<i>AMM → Buyback</i>	Quasi-arbitrage opportunity implied by cross-mechanism price signals and defined as an economically feasible ordered pair of buying a RWA token via the AMM and selling via Buyback. Importantly, these opportunities are inferred from price signals and need not be feasible at economically meaningful scale or implementable for all investors due to execution risk, quantity caps, and whitelisting.	Own
<i>P2P ask → AMM</i>	Quasi-arbitrage opportunity implied by cross-mechanism price signals and defined as an economically feasible ordered pair of buying a RWA token via the P2P ask and selling via the AMM. Importantly, these opportunities are inferred from price signals and need not be feasible at economically meaningful scale or implementable for all investors due to execution risk, quantity caps, and whitelisting.	Own
<i>P2P ask → P2P bid</i>	Quasi-arbitrage opportunity implied by within-mechanism P2P offer-book price signals (ask vs bid) and defined as an economically feasible ordered pair of buying a RWA token via the P2P ask and selling into the P2P bid. Importantly, these opportunities are inferred from price signals and need not be feasible at economically meaningful scale or implementable for all investors due to execution risk, quantity caps, and whitelisting.	Own
<i>P2P ask → Buyback</i>	Quasi-arbitrage opportunity implied by cross-mechanism price signals and defined as an economically feasible ordered pair of buying a RWA token via the P2P ask and selling via Buyback. Importantly, these opportunities are inferred from price signals and need not be feasible at economically meaningful scale or implementable for all investors due to execution risk, quantity caps, and whitelisting.	Own
<i>P2P ask ⇒ AMM</i>	Realized round-trip quasi-arbitrage sequence defined as a buy via the P2P ask followed by a sell via the AMM in the same RWA token within a fixed window, with strictly positive net profit after fees.	Own
<i>P2P ask ⇒ P2P bid</i>	Realized round-trip quasi-arbitrage sequence defined as a buy via the P2P ask followed by a sell into the P2P bid in the same RWA token within a fixed window, with strictly positive net profit after fees.	Own
<i>P2P ask ⇒ Buyback</i>	Realized round-trip quasi-arbitrage sequence defined as a buy via the P2P ask followed by a sell via Buyback in the same RWA token within a fixed window, with strictly positive net profit after fees.	Own
Dependent variables (intraday perspective)		
<i>AMM trade</i>	Indicator equal to one if the trade executes via AMM.	Own
<i>P2P trade</i>	Indicator equal to one if the trade executes via P2P.	Own
<i>Buyback trade</i>	Indicator equal to one if the trade executes via Buyback.	Own
Dependent variables (daily perspective)		
<i>External liquidity (USD)</i>	USD value of the AMM pool liquidity not attributable to the issuer-designated LP address (i.e., provided by non-issuer LPs).	Own
Stress event variables—intraday perspective (± 100 -days around the shock)		
<i>1-token price impact (bps)</i>	Implied execution cost, in basis points, of a hypothetical 1-token trade (\$50). For AMMs, computed from contemporaneous pool states under constant-product pricing as $10,000 \times \left(1 - \frac{X}{X+1}\right)$, where $X = L^{\text{pool}}/P^{\text{spot}}$ denotes the implied base-token reserve inferred from total pool value and the AMM spot price. For P2P trades, defined as $10,000 \times (P^{\text{exec}} - P^{\text{ask}})/P^{\text{ask}}$, using the contemporaneous liquidity-weighted ask quote.	Own
<i>P2P bid–ask spread (bps)</i>	Implied cost of immediacy in peer-to-peer trading, expressed in basis points. Computed from contemporaneous bilateral quotes as $10,000 \times (P^{\text{ask}} - P^{\text{bid}})/P^{\text{mid}}$, where $P^{\text{mid}} = \frac{1}{2}(P^{\text{ask}} + P^{\text{bid}})$ and P^{ask} and P^{bid} denote the liquidity-weighted ask and bid quotes prevailing at the time of observation. The measure captures trading costs conditional on two-sided P2P liquidity.	Own
<i>Asset size (USD)</i>	Concurrent USD appraisal value of the underlying real estate investment at the time of trade, measured by the total investment value associated with the traded token. The variable reflects the contemporaneous fair value used in issuer Buyback and serves as a proxy for underlying asset size. We use the token-level median of this measure to partition the sample into small and large assets in cross-sectional analyses.	Own
<i>TRO affected</i>	Indicator variable equal to one if a trade involves a RWA token directly affected by the Temporary Restraining Order (TRO), and zero otherwise. On July 22, 2025, a Detroit judge granted the City of Detroit a TRO requiring rental collections on affected properties to be halted until certificates of compliance are obtained. As a result, rental income distributions to investors were halted and redirected to escrow accounts, and the Buyback was suspended by the issuer for all affected properties.	Own
<i>Total fees (USD, <i>wavg</i>)</i>	Token–day average of <i>Total fees (USD)</i> across all trades in that token on that day, weighted by <i>Trade volume (USD)</i> .	Own
Market control variables (daily perspective)		

<i>RWA token price volatility (7-day)</i>	7-day rolling standard deviation of daily volume-weighted average prices for a given RWA token.	Own
<i>Market-to-appraisal ratio</i>	Ratio of total RWA market capitalization based on 7-day rolling volume-weighted average prices to total RWA market capitalization based on concurrent Buyback prices.	Own
<i>Cumulative return ETH (one week before)</i>	Cumulative return of Ether over the prior week.	Coinbase
<i>Average ETH tx fee (USD)</i>	Average Ethereum transaction fee over the observation period, expressed in USD.	Coinbase
<i>1m Treasury yield Δ</i>	Daily log-change of the market yield on US Treasury securities at 1-month constant maturity (investment basis).	FRED
<i>10y Treasury yield Δ</i>	Daily log-change of the market yield on US Treasury securities at 10-year constant maturity (investment basis).	FRED
<i>ADS Index Δ</i>	Daily log-change of the Aruoba–Diebold–Scotti (ADS) Business Condition Index (Aruoba et al., 2009).	Fed Philadelphia
<i>S&P Case-Shiller Index Δ</i>	Monthly log-change of the US S&P Case-Shiller National Home Price Index, lagged by one month.	FRED
Trader heterogeneity		
<i>All trader wallets</i>	Seller wallets with at least one observed secondary-market trade in the sample.	Own

Note: This table lists all variables, their definitions, and data sources. *Own* denotes authors’ own calculations based on on-chain transaction records.

A.4 Data in detail

This appendix summarizes the construction and processing of our on-chain datasets for the RealT’s tokenized RWA market. All steps are implemented in Python and R. The full codebase, APIs, and intermediate data outputs are included in the replication package.

A.4.1 Data description

Our primary dataset is a complete transaction-level panel of all secondary-market trades in RealT-issued RWA tokens from Nov. 4, 2019 to Oct. 16, 2025, covering AMM, P2P, and Buyback activity. It contains 777,038 trades (\$41.2 million) across 9,202 unique wallets.¹¹ Summary statistics are reported in Tables 2.1 and 2.2. We merge trades with token- and property-level fundamentals for 752 tokenized RWAs: 728 are equity claims on SPVs holding one or multiple residential rental properties in the United States and Latin America, and the remainder are property-loan tokens, factoring tokens, and tokenized company shares. Overall, the sample covers 923 properties and 1,794 residential units.

A.4.2 Data sources and extraction

Our dataset integrates on-chain event data stored on Ethereum mainnet and Gnosis Chain, whose sources are explained in this subsection.

Identification of relevant smart-contract addresses. We begin by identifying the smart-contract addresses that mediate trading on each marketplace. Starting from the RWA

¹¹Multiple wallets may belong to the same individual, implying potential wash trading (Cong et al., 2023).

token contract addresses shown on RealT’s primary-market listing pages (RealT Website), we retrace transactions and counterparties using the respective blockchain explorers, Etherscan for Ethereum mainnet and Gnosisscan for Gnosis Chain. To further validate contract identification, one author executed small test trades (below \$100) on each venue and retraced the resulting on-chain interactions. These steps yield the recurring *interaction-contract addresses*, i.e., the smart-contract addresses that traders’ transactions call when executing trades on a given venue, listed below:

- **P2P.**

- Swapcat: 0xB18713Ac02Fc2090c0447e539524a5c76f327a3b (Ethereum and Gnosis)
- YAM: 0xc759aa7f9dd9720a1502c104dae4f9852bb17c14 (Ethereum and Gnosis)

- **AMM.**

- Levinswap: 0xb18d4f69627F8320619A696202Ad2C430CeF7C53 (Gnosis only)
- Uniswap pools do not have a single protocol-level interaction-contract address; events are emitted by the pool (pair) contracts.

- **Buyback.** Buybacks execute directly against the issuer-controlled, token-specific vault (property vault), not a centralized protocol contract. We identify each vault by retracing the initial issuance flow: the earliest transfers originate from the zero address (0x0...0), and the first recipient that subsequently distributes tokens in systematic primary-market sales is the vault. We then classify transfers by direction relative to the vault: OUT (vault→investors) as primary-market sales and IN (investors→vault) as issuer repurchases (Buybacks). We validate this mapping by reconciling cumulative OUT flows with issued supply, accounting for RealT’s standard issuer retention (5%) and excluding rare cases where the issuer supplies inventory directly to secondary venues (STO-like sales). As a further consistency check, we use RealT’s main operating wallet (0x5Fc96c182Bb7E0413c08e8e03e9d7EFc6cf0B099) to confirm that the identified vault engages in the expected issuance and cash-management interactions.

These central addresses serve as fixed reference points throughout our data extraction pipeline and are used for manual cross-checks at each step (contract identification, event retrieval, and test reconstruction outputs using alternative methods involving Etherscan and Gnosisscan APIs).

Indexed on-chain data. To retrieve transaction- and state-level data at scale, we query official RealT and community-maintained subgraphs on The Graph, which index raw blockchain logs into structured entities. The subgraphs used in this paper are maintained under a common creator address 0x1bcfe8666cbb7edd3eda5d343d5f2d4ce853b034. We access each subgraph listed below via The Graph using the endpoint [https://gateway.thegraph.com/api/\[api-key\]/subgraphs/id/\[subgraph-id\]](https://gateway.thegraph.com/api/[api-key]/subgraphs/id/[subgraph-id]). An API key can be created free of charge by connecting a wallet to The Graph, and the free-tier rate limits are sufficient to reproduce our full extraction.

Marketplace and subgraph IDs (by blockchain).

- **P2P**

- Swapcat: 4zYp9FYpdepqxqgMCwDTawDvfFbS9VsXmocyiSFLwEvi7 (Ethereum)
- Swapcat: BC5uHuHf4YybNfQECjvHJyZmEfLrPSUu8fLNMZsQ38jx (Gnosis)
- YAM: gMUyhime7v5vkmCv8uJV3L8L3MnLqx2BjTjtiBh6zhk (Ethereum)
- YAM: 2RHaLwFYozfHN2bbwUmAoQeJuGn7FX5UREGckZFwfiZi (Gnosis)

- **AMM**

- Levinswap: 6fus6wPYtbNAXr5ZwTrP7anuPypFkEw4snyWFXWbSbAj (Gnosis)
- Uniswap V2: 2zkYvVc8xsmytcEjE9pr523qFcbvqRGxhuGGtraozs66 (Ethereum)
- Uniswap V1: ESnjgAG9NjfmHypk4Huu4PVvz55fUwpyrRqHF21thoLJ (Ethereum, not RealT-curated)

Data validation. We validate indexed records in multiple ways. First, we manually cross-check selected indexed transactions against multiple blockchain explorer outputs, including Etherscan, Gnosisscan, and Blockscout (as an alternative for both) and compare and recompute key decoding logic (e.g., offer enumerations for P2P offer-books, AMM pool-reserve transitions, and implied pool prices). Second, we cross-check subgraph indexer consensus using The Graph’s Query Cross Checking Tool and rerun identical queries to verify stability across indexers and across repeated calls. Third, at the time of collection, the relevant subgraphs were served by multiple indexers, mitigating single-indexer failures. Finally, for AMMs, we independently recompute pool reserve transitions from swap events and retrace selected liquidity-provider balances and fee accrual logic for spot checks. We further cross-check our results in aggregate with alternative analytical and smart-contract interface applications for RealT’s RWA tokens (RealT Lab) and RealT as a platform itself (e.g. DefiLlama).

Further data extraction methods. For the Buyback transactions we use the Etherscan and Gnosisscan APIs for Gnosis and Ethereum and fetch all associated ERC-20 transfers with the corresponding property vault addresses.

A.4.3 Event-level construction and enrichment

Event retrieval and schema. Using the subgraphs listed above, we query all marketplace events at the block level. For each event, indexed records include the event type (e.g., `mint`, `swap`, `burn` for AMMs; `created`, `updated`, `purchase`, `removed` for P2P), transaction hash, log index, block number, timestamp, sender and receiver addresses, token symbols and addresses, and traded quantities. Mechanism-specific fields include pool identifiers, LP fee parameters, LP token holder identifiers and positions (AMMs), and offer identifiers and offer-state variables (P2P).

State reconstruction. For P2P, the event history allows reconstruction of the full offer-book dynamics by offer ID, including updates, partial fills, and removals. For AMMs, we reconstruct liquidity-pool state trajectories, including reserve transitions and implied prices, as well as liquidity provision and withdrawal activity and the corresponding LP identities at the pool level.

AMM liquidity provision and external participation. We enrich our AMM event data with on-chain liquidity information from three The Graph subgraphs mentioned above: Uniswap V1 (Ethereum), Uniswap V2 (Ethereum), and LevinSwap (Gnosis Chain). Liquidity providers (LPs) receive pool-specific LP tokens when adding liquidity and surrender them when withdrawing liquidity. These LP tokens represent pro-rata claims on pool reserves and trading fees. We separate issuer-provided from externally provided liquidity by defining issuer-related addresses:

- **Issuer-related LP addresses.** We define issuer-related addresses based on (i) repeated pool seeding (earliest `mint` recipient per pool), (ii) involvement in protocol/administrative and buyback-related flows, and (iii) persistent linkages in LP-token transfers around the Ethereum→Gnosis migration 2021. The issuer-related set mainly includes:
 - RealT main pool seeder (> 95% of RWA pools):
 - * 0x5fc96c182bb7e0413c08e8e03e9d7efc6cf0b099
 - Exclusive Uniswap V1 seeder: 0x387ca78a5edf943d267b9248cce1b0fdb67d0029
 - * Ethereum-Gnosis-migration-linked interim addresses (LP-token transfer linkages):
 - * 0xcba15d7eb8fe219a70df21d3259c3cbbf7c6e525,
 - * 0xf22c2e3475e4a066f4e9f44567c950dd36112d05
 - Protocol operational contract (*LevinMaker*):
 - * 0x6d81dda24b7ff5b4a65039ff15d06a076e018e49

Issuer vs. external liquidity provision. For each pool, we retrieve total LP token supply and LP token balances held by issuer-related addresses. We query these at the pool’s first observed block and at all `mint`/`burn` blocks, then forward-fill within pool between queried blocks. Issuer liquidity is measured by issuer-held LP tokens; external liquidity is the remainder (total supply minus issuer-held LP tokens). We measure *net liquidity provision* as signed changes in LP token balances at `mint`/`burn` blocks (positive deposits, negative withdrawals), aggregated by pool-day and split by issuer-related vs. external LPs. Using daily end-of-day snapshots of LP token holders (positive balances) excluding issuer-related addresses, we further report the number of active external LPs and the cumulative number of unique external LPs over time.

Transaction fees and USD conversions. Using transaction hashes, we augment on-chain events with realized network fees from the Etherscan and Gnosisscan APIs. USD conversions use the Coinbase API: ETH-denominated values are mapped to USD using hourly ETH/USD rates. For LEVIN, we infer LEVIN/USD from contemporaneous WXDAI/LEVIN pool reserves via the implied spot price (WXDAI being USD-pegged on Gnosis).

Property-level and token fundamentals We collect token-, asset-, and platform-level fundamentals from the RealT website, the community-maintained RealT API, and RealT’s social media communications. We use token metadata, STO terms, rental-income histories, and appraisals to enrich the on-chain panels with asset characteristics and platform events. Time-invariant attributes include issuance volume, token supply, issuance date, location, and

property characteristics. Time-varying attributes include rental income, rent-to-value yields, rental status, building age, and appraisal values.

A.4.4 Sample filters

Our sample spans Nov. 4, 2019, 16:07:31 (Ethereum block 8872285) to Oct. 16, 2025, 08:59:35 (Gnosis block 42651526). Block timestamps are in UTC. To ensure mechanical precision and remove economically implausible records, we further apply the following mechanism-specific rule-based filters.

AMM. We exclude four AMM swaps with implausible implied prices above \$1,000,000 and one swap below \$0.10. These bounds are intentionally conservative to avoid over-filtering in thin pools where deterministic AMM pricing can produce large dispersion. We do not winsorize or otherwise smooth AMM price signals, since extreme quotes are informative about low depth and are reported as such (see Table 2.2). From AMM subgraphs, we reconstruct liquidity-pool states (reserves, implied prices, execution prices, slippage, and liquidity-provider positions), yielding a pool-level panel.

P2P. We restrict attention to offers and purchases that involve at least one RWA token and are quoted against established stablecoins (USDC, USDT, WXDAI, DAI, REUSD, MAI). We drop all other pairs, including RWA token–RWA token transfers, which removes 636 offers and 349 purchases. Because wallets can post offers exceeding their on-chain balances, we discard offers whose advertised quantity exceeds 30% of total token supply. This threshold is above the 15% primary-market allocation limit and remains permissive of secondary-market accumulation. This removes 10,757 offers and 1,108 purchases. We further exclude offers hard-deleted by RealT administrators during three documented purge blocks on YAM (619 offers and 1,191 purchases) and observations with missing or implausible USD prices outside the \$1–\$1,000 range (1,377 offers and 307 purchases). Together, these filters remove approximately 6.9% of offer events and 1.4% of purchase events, leaving 179,788 offers and 207,885 purchases. Purchases exceed offers because a single offer can be partially filled multiple times, generating multiple purchase records.

Buyback. One Buyback transaction is identified as an internal transfer and excluded; all remaining observations are retained. Buyback transactions are identified by tracing token flows between issuer and external wallets and netting these flows against STO issuance volumes.

A.4.5 Reproducibility

For reproducibility, our replication package includes scripts to rerun the full extraction and reconstruction pipeline. Because 99% of secondary-market trading volume in our sample settles on Gnosis Chain, we run the Gnosis portion against a local archive node to ensure access to historical state at arbitrary block heights and to avoid reliance on third-party indexer availability. The archive node is queried via a Chainstack Remote Procedure Call (RPC) endpoint, and all intermediate outputs are written to disk to permit exact reruns and step-wise verification.

A.5 Quasi-arbitrage analysis

Quasi-arbitrage. To reflect market frictions, including execution uncertainty, heterogeneous mechanisms, and appraisal-anchored pricing, we term the exploitation of cross-mechanism price dispersion *quasi-arbitrage* rather than risk-free arbitrage. Panel A of Table A.6 shows that median prices are tightly aligned across mechanisms, while the mean AMM quote and its dispersion are substantially larger, reflecting a small number of very thin pools. Panel B translates these price signal differentials into token-day quasi-arbitrage opportunities for each economically feasible ordered pair of buy and sell mechanisms, denoted as $m^{\text{buy}} \rightarrow m^{\text{sell}}$. Importantly, these opportunities are inferred from contemporaneous price signals rather than executed trades and therefore need not be feasible at economically meaningful scale, nor implementable for all investors given quantity caps, execution risk, and whitelist-based access restrictions. Panel C reports realized round-trip quasi-arbitrage buy-sell sequences, defined as a buy followed by a sell in the same token within a fixed window (1 minute or 1 hour) with strictly positive net profit after fees. We denote such sequences by $m^{\text{buy}} \Rightarrow m^{\text{sell}}$. To capture cross-mechanism price pressure even when trades route through intermediaries or smart-contract routers, Panel C imposes no wallet identity constraint, yielding conservative lower-bound estimates of actual trading activity. Across both horizons, realized quasi-arbitrage is concentrated in *P2P ask $\Rightarrow$ AMM* and *P2P ask $\Rightarrow$ P2P bid* sequences.

Table A.6: Price signals and quasi-arbitrage

<i>Panel A: Summary statistics of price signals</i>				
Mechanism	Token-days	Median (USD)	Mean (USD)	St. Dev. (USD)
Price signal AMM (USD)	187,749	52.67	62.44	5,013.91
Price signal P2P bid offer (USD)	79,931	49.64	51.39	13.46
Price signal P2P ask offer (USD)	158,232	55.00	56.66	11.19
Price signal Buyback (USD)	188,198	50.79	51.14	13.99
<i>Panel B: Quasi-arbitrage opportunities from price signals (token-day level)</i>				
Buy $\rightarrow$ Sell	Token-days with opportunity	Median spread (USD)	Mean spread (bps)	Median spread (bps)
AMM $\rightarrow$ P2P bid	7,546	1.34	567.60	258.29
AMM $\rightarrow$ Buyback	48,426	1.13	385.14	217.91
P2P ask $\rightarrow$ AMM	45,350	2.09	757.18	362.98
P2P ask $\rightarrow$ P2P bid	1,897	1.22	418.73	233.39
P2P ask $\rightarrow$ Buyback	8,567	0.50	285.42	95.69
<i>Panel C: Realized quasi-arbitrage trades (transaction-pair level, no wallet constraint)</i>				
Buy $\Rightarrow$ Sell	Token-pairs	Token-quantity	Net-profit (USD)	Net-volume (USD)
<i>1-minute window</i>				
P2P ask $\Rightarrow$ AMM	10,399	1,302.12	4,305.86	80,743.22
P2P ask $\Rightarrow$ P2P bid	1,752	2,792.58	2,591.28	150,016.17
<i>1-hour window</i>				
P2P ask $\Rightarrow$ P2P bid	21,573	41,692.72	106,187.95	2,279,314.96
P2P ask $\Rightarrow$ AMM	63,590	3,489.56	11,215.12	204,863.33
P2P ask $\Rightarrow$ Buyback	42	46.02	312.07	3,088.01

Note: Panel A reports token-day coverage (Token-days) and the median, mean, and standard deviation computed across transaction-level signal observations on those token-days. Panel B reports token-day quasi-arbitrage opportunities implied by these signals and denoted by $m^{\text{buy}} \rightarrow m^{\text{sell}}$ (buy via m^{buy} , sell via m^{sell}). Panel C reports realized quasi-arbitrage trades denoted by $m^{\text{buy}} \Rightarrow m^{\text{sell}}$ and defined as a buy followed by a sell in the same token within the indicated time window, with strictly positive net profit after fees; net volume is the sell-leg notional at net prices. Please note that, especially for P2P, many realized sequential trades occur within a single day, relative to the token-day level of analysis in Panel B. The AMM price signal is identical for bid and ask by construction, so we do not distinguish between them in our notation. The Buyback mechanism provides only a bid quote, as tokens can only be sold into it.

References

- Adams H., Zinsmeister N., and Robinson D. H. (2020). Uniswap v2 core. URL: <https://app.uniswap.org/whitepaper.pdf>.
- Almgren R. and Chriss N. (2001). Optimal execution of portfolio transactions. *Journal of Risk* 3.2, pp. 5–39.
- Angeris G., Kao H.-T., Chiang R., Noyes C. R., and Chitra T. (2019). An analysis of Uniswap markets. *Cryptoeconomic Systems Journal* 1.1.
- Aoyagi J. and Ito Y. (2025). Coexisting exchange platforms: Limit order books and automated market makers. *Journal of Political Economy Microeconomics* 3.3, pp. 611–648.
- Aquilina M., Foley S., Gambacorta L., and Krekel W. (2024). Decentralised dealers? Examining liquidity provision in decentralised exchanges. BIS Working Papers 1227. Bank for International Settlements.
- Aruoba S. B., Diebold F. X., and Scotti C. (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics* 27.4, pp. 417–427.
- Barbon A. and Ranaldo A. (2026). On the quality of cryptocurrency markets: Centralized vs. decentralized exchanges. *Management Science. Articles in Advance*.
- Bergkamp M., Sifat I., and Swinkels L. (2025). Market maturation and democratization effects of tokenized real estate. SSRN. Available at URL: <https://ssrn.com/abstract=5195172>.
- Brauneis A., Mestel R., Riordan R., and Theissen E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking and Finance* 124, p. 106041.
- Buterin V. (2014). Ethereum: A next-generation smart contract and decentralized application platform. White Paper. URL: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- Capponi A. and Jia R. (2021). The adoption of blockchain-based decentralized exchanges. arXiv. Available at URL: <https://arxiv.org/abs/2103.08842>.
- Capponi A. and Jia R. (2025). Liquidity provision on blockchain-based decentralized exchanges. *The Review of Financial Studies* 38.10, pp. 3040–3085.
- Capponi A., Jia R., and Yu S. (2026). Price discovery on decentralized exchanges. *The Review of Financial Studies*, hhag002.
- Capponi A., Jia R., and Zhu B. (2024). The paradox of just-in-time liquidity in decentralized exchanges: More providers can lead to less liquidity. arXiv. Available at URL: <https://arxiv.org/abs/2311.18164>.
- Case K. E. and Shiller R. J. (1989). The efficiency of the market for single-family homes. *American Economic Review* 79.1, pp. 125–137.
- Chen D. and Duffie D. (2021). Market fragmentation. *American Economic Review* 111.7, pp. 2247–2274.
- Cong L. W. and He Z. (2019). Blockchain disruption and smart contracts. *The Review of Financial Studies* 32.5, pp. 1754–1797.
- Cong L. W., Landsman W. R., Rabetti D., Zhang C., and Zhao W. (2025a). Tokenized stocks. SSRN. Available at URL: <https://ssrn.com/abstract=5937314>.

- Cong L. W., Li X., Tang K., and Yang Y. (2023). Crypto wash trading. *Management Science* 69.11, pp. 6427–6454.
- Cong L. W., Mayer S., and Rabetti D. (2025b). Tokenizing real-world assets. Working Paper.
- Cong L. W., Prasad E. S., and Rabetti D. (2025c). Financial and informational integration through oracle networks. Working Paper 33639. National Bureau of Economic Research.
- Duffie D., Gârleanu N., and Pedersen L. H. (2005). Over-the-counter markets. *Econometrica* 73.6, pp. 1815–1847.
- Goetzmann W. N., Spaenjers C., and Van Nieuwerburgh S. (2021). Real and private-value assets. *The Review of Financial Studies* 34.8, pp. 3497–3526.
- Harvey C. R., John K., and Saleh F. (2025). Productivity enables security: The economics of blockchain settlement. SSRN. Available at URL: <https://ssrn.com/abstract=5868863>.
- Harvey C. R. and Rabetti D. (2024). International business and decentralized finance. *Journal of International Business Studies* 55.7, pp. 840–863.
- Hasbrouck J. (1995). One security, many markets: Determining the contributions to price discovery. *The Journal of Finance* 50.4, pp. 1175–1199.
- Heimbach L., Wang Y., and Wattenhofer R. (2021). Behavior of liquidity providers in decentralized exchanges. arXiv. Available at URL: <https://arxiv.org/abs/2105.13822>.
- Kreppmeier J., Laschinger R., Steininger B. I., and Dorfleitner G. (2023). Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *Journal of Banking and Finance* 154, p. 106940.
- Kyle A. S. (1985). Continuous auctions and insider trading. *Econometrica* 53.6, pp. 1315–1335.
- Lambert T., Liebau D., and Roosenboom P. (2022). Security token offerings. *Small Business Economics* 59, pp. 299–325.
- Lee S., Lee J., and Lee Y. (2023). Dissecting the Terra-LUNA crash: Evidence from the spillover effect and information flow. *Finance Research Letters* 53, p. 103590.
- Lehar A. and Parlour C. A. (2025). Decentralized exchange: The Uniswap automated market maker. *The Journal of Finance* 80.1, pp. 321–374.
- Lehar A., Parlour C. A., and Zoican M. (2024). Fragmentation and optimal liquidity supply on decentralized exchanges. SSRN. Available at URL: <https://ssrn.com/abstract=4267429>.
- Li X., Ke R., Singh S. F., Park A., and Veneris A. (2024). On tokenizing securities in contemporary decentralized finance ecosystems. 2024 6th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS). Berlin, Germany: IEEE, pp. 1–9.
- Makarov I. and Schoar A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics* 135.2, pp. 293–319.
- Malamud S. and Rostek M. (2017). Decentralized exchange. *American Economic Review* 107.11, pp. 3320–3362.
- Malinova K. and Park A. (2017). Market design with blockchain technology. SSRN. Available at URL: <https://ssrn.com/abstract=2785626>.
- Milionis J., Moallemi C. C., Roughgarden T., and Zhang A. L. (2024). Automated market making and loss-versus-rebalancing. arXiv. Available at URL: <https://arxiv.org/abs/2208.06046>.

- O'Hara M. (2003). Presidential address: Liquidity and price discovery. *The Journal of Finance* 58.4, pp. 1335–1354.
- Pagano M. (1989). Trading volume and asset liquidity. *The Quarterly Journal of Economics* 104.2, pp. 255–274.
- Park A. (2023). The conceptual flaws of decentralized automated market making. *Management Science* 69.11, pp. 6731–6751.
- Park A. and Stinner J. (2025). Phantom liquidity in decentralized lending. SSRN. Available at URL: <https://ssrn.com/abstract=4489732>.
- Rostek M. and Yoon J. H. (2025). Financial product design in decentralized markets. *Journal of Political Economy* 133.3, pp. 888–934.
- Schär F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review* 103.2, pp. 153–174.
- Swinkels L. (2023). Empirical evidence on the ownership and liquidity of real estate tokens. *Financial Innovation* 9, pp. 1–29.
- Uhlig H. (2022). A Luna-tic stablecoin crash. Working Paper 30256. National Bureau of Economic Research.
- Whitaker A. and Kräussl R. (2020). Fractional equity, blockchain, and the future of creative work. *Management Science* 66.10, pp. 4594–4611.
- Williams J. T. (1995). Pricing real assets with costly search. *The Review of Financial Studies* 8.1, pp. 55–90.
- Zhu B. Z., Liu D., Wan X., Liao G., Moallemi C. C., and Bachu B. (2025). What drives liquidity on decentralized exchanges? Evidence from the Uniswap protocol. arXiv. Available at URL: <https://arxiv.org/abs/2410.19107>.

Conflicts of Interest

One author bought and sold few digital tokens issued by the company RealT so that they could retrace the process and trading of tokenization. The current value of these holdings is less than \$100.

Acknowledgments

We explicitly thank Lin William Cong, Fahad Saleh, Kose John, Tao Li, and the other participants at the American Finance Association Annual Meeting 2026. We also explicitly thank Agostino Capponi, Maureen O'Hara, Peter O'Neill, Sean Foley, and the participants at the 7th Sydney Market Microstructure Conference; Tong Yu and the participants at the 38th Australasian Banking and Finance Conference; Siqu Zheng, James Scott, and the participants at the 2025 MIT Symposium on Technology, AI, and Real Estate; Bertram Steininger and the participants at the 17th Real Estate Markets and Capital Markets Network (ReCapNet) conference; Lennart Sperling and the participants at the 31st Annual Meeting of the German Finance Association (DGF); and Pierre Collin-Dufresne and the participants at Munich Finance Day. We also appreciate the helpful comments and suggestions from William N. Goetzmann and the participants at the 2025 American Real Estate Society Annual Meeting and Conference in Tucson, USA, as well as Ryan Riordan, Ralf Elsas, and the participants at the Workshop European Finance and Microstructure Revisited. Furthermore, we thank Julian Prat, Michele Fabi, Andrew Macpherson, Engin Iyidogan, Hamza El Khalloufi, and the participants at the Tokenomics Workshop 2024 in Tokyo, Japan. These comments have been invaluable for improving our work. All remaining errors are our own.

Chapter 3

Tokenization, Transparency, and Time on Market: Empirical Evidence from a Blockchain-Based Real Estate Marketplace

This paper is co-authored with Ralf Laschinger (Institute for Financial Innovation and Technology, LMU Munich School of Management) and Wolfgang Schäfers (IREBS, International Real Estate Business School, University of Regensburg), and has received a minor revision decision for the *Journal of Real Estate Research* Special Issue on AI and Real Estate.

Abstract: Tokenized real-estate markets, where fractional digital claims on individual properties are traded, create an unusually transparent setting in which listing histories, price revisions, and execution outcomes are publicly observable. This paper applies standard Time on Market (ToM) frameworks to tokenized real estate for the first time and studies listing behavior, seller adjustment, and marketing time for individual real-estate tokens. Using complete listing and transaction data from a large blockchain-based peer-to-peer marketplace, we compile a unique dataset covering more than 111,000 listings and more than 72,000 sales for 731 tokenized rental residential properties in the United States and Latin America from 2020 to 2025. Using OLS and accelerated failure time models, we find that overpricing relative to the redemption benchmark prolongs ToM, whereas each additional percentage point of downward revision shortens remaining ToM by about 6%, conditional on the premium remaining after the revision. Holding total price concessions constant, listings sold after a single decisive price reduction achieve higher final premia than listings repriced gradually. Overall, the study shows that tokenization makes liquidity and execution outcomes more tightly linked to transparent relative pricing and publicly observable seller adjustment, with implications for increasingly transparent digital listing platforms in real-estate markets more broadly.

Keywords: Tokenization, Time on Market, Price Revisions, Liquidity, Blockchain

JEL: G12, G14, O33, R30

3.1 Introduction

Real estate markets are characterized by illiquidity, asset heterogeneity, decentralization, and limited transparency over the selling process. These frictions complicate search and matching and contribute to uncertain market outcomes with respect to pricing and time on market. In response to these frictions, digital technologies are increasingly reshaping real-estate intermediation and market outcomes, for example through online marketplaces, brokerage technologies, and tools such as automated valuation models, instant-buying options, or virtual tours. More recently, blockchain-based tokenization has emerged as a new technology to potentially reshape the transaction process and market microstructure of real estate exchange. More broadly, tokenization refers to representing asset-related claims as small, transferable digital units recorded on a blockchain and has attracted growing attention in research and practice, including from large asset managers and financial institutions.¹ In real estate, tokenization aims above all to make a lumpy and illiquid asset more divisible, more tradable, and easier to transfer. It moves traditionally opaque real estate transaction and listing histories into a digitally recorded environment in which listing, repricing, and execution behavior can be observed far more closely than in conventional property markets. In doing so, it allows classic questions on overpricing and price adjustment in the Time on Market (ToM) literature to be revisited in a new, more transparent technological market setting (Haurin, 1988; Anglin et al., 2003; Herrin et al., 2004).

In this paper, we study a peer-to-peer marketplace for tokenized real estate, where rental properties are fractionalized into blockchain-based tokens that can be listed and traded individually. Unlike conventional real estate datasets, which typically observe the initial list price and the transaction price but not the intermediate sequence of revisions and removals, this setting records the full listing history from initial posting through price changes to sale or removal (Knight, 2002). The market operates through posted, take-it-or-leave-it listings rather than bilateral bargaining, placing the listing problem closer to posted-price financial markets than to negotiated property sales. At the same time, it retains key real-estate anchors such as heterogeneous property characteristics, locational effects, rental income streams, and periodic valuation benchmarks. Using complete listing and transaction data for more than 111,000 listings, more than 72,000 sales, and 731 tokenized rental residential properties in the United States and Latin America from 2020 to 2025, and employing OLS and accelerated failure time models, we study ToM dynamics, price revision strategies, and selling outcomes at a granularity unavailable in conventional property markets.

The aim of this study is to examine how traceable listing histories shape ToM, seller adjustment, and execution outcomes in a tokenized real-estate marketplace. Although a large real-estate literature studies listing prices, ToM, and price revisions, much of this evidence comes from conventional brokered markets, where bargaining weakens the link between posted and transaction prices and intermediate revisions are often unobserved, so that the observed listing path remains only an imperfect proxy for the economically relevant price path. Our setting allows us to revisit these classic questions in a more transparent environment by combining publicly traceable listing histories with take-it-or-leave-it execution and a continuously observable valuation benchmark. At the same time, existing empirical work on

¹For example, in his 2025 Annual Chairman’s Letter, BlackRock CEO Larry Fink presented tokenization as a development with the potential to fundamentally alter how financial and private-market assets, including real estate, are held, transferred, and settled (BlackRock, 2025).

tokenized real-estate markets has focused primarily on issuance activity, ownership structure, and liquidity organization rather than on listing behavior, ToM, and seller adjustment at the level of individual token listings (Swinkels, 2023; Kreppmeier et al., 2023; Bergkamp et al., 2025; Laschinger et al., 2026). Against this background, we ask two questions: first, what determines ToM in a tokenized real-estate marketplace; and second, how do price revision strategies affect remaining ToM and price outcomes when the full listing path is publicly observable?

We find that overpricing relative to the redemption benchmark strongly prolongs ToM, confirming search-theoretic predictions in a setting with a continuously observable outside option. Larger downward price revisions are associated with shorter remaining ToM: each additional percentage point of downward revision is associated with roughly 6% shorter remaining duration, conditional on the degree of overpricing that remains after the revision. Holding the total price concession constant, listings that implement the adjustment in a single cut achieve higher final selling premia than otherwise similar listings that spread the same concession across multiple reductions, consistent with repeated, publicly observable revisions signaling a declining reservation price and inviting buyer delay.

The results show that in a market where the full listing history is public, both the magnitude and the path of price adjustment are first-order determinants of liquidity and execution outcomes. The paper makes three contributions. First, it provides the first systematic evidence on ToM, listing behavior, and seller adjustment for individual tokenized real-estate listings. Second, it re-examines classic questions on overpricing, marketing time, and seller adjustment in a setting where the posted price is the execution price and the full listing path is observable. These insights may also extend to traditional real-estate listing platforms, which make offer histories increasingly visible to buyers. Third, it shows that the economic importance of tokenization lies not only in divisibility, transferability, or lower trading frictions, but also in moving real-estate exchange into a more transparent and traceable microstructure in which relative pricing, market thickness, and publicly observable seller adjustments become more consequential for liquidity and execution outcomes.

The remainder of the paper is organized as follows. Section 3.2 provides background on real-estate tokenization and the institutional marketplace setting. Section 3.3 reviews the related literature and develops the hypotheses. Section 3.4 describes the data, and Section 3.5 presents the methodology. Section 3.6 reports the results, and Section 3.7 concludes.

3.2 Background on real estate tokenization and institutional marketplace details

3.2.1 Real estate tokenization

Real estate tokenization represents ownership or cash-flow claims on real estate using digital tokens recorded on blockchains. Tokenization can take several forms, including whole-property tokenization, fractionalized ownership claims, or the tokenization of specific property cash flows. We focus on the form currently most relevant for active secondary trading: fractional-

ized, transferable claims linked to property-holding vehicles or their cash flows. Tokenization in this form combines fractionalization and splitting exposure into small, transferable units, with blockchain-based transfer and settlement infrastructure that can automate payments and make transactions verifiable and traceable. This lowers participation frictions through smaller ticket sizes, increased tradability, faster settlement, and potentially broader investor access (Benedetti and Rodríguez-Garnica, 2023). Industry forecasts project that tokenization of illiquid assets could reach \$16 trillion by 2030, with tokenized real estate potentially reaching \$4 trillion by 2035 (BCG, 2022; Deloitte, 2025). These figures, while not yet realized, reflect growing interest among market participants in blockchain-based real estate investment infrastructure.

Within the real estate investment landscape, fractionalized tokenized exposures sit between direct property ownership and pooled vehicles (Liu and Chen, 2025). Unlike REITs or open-end funds, which provide portfolio exposure with limited discretion over individual assets, most real estate token offerings are currently linked to single properties, enabling more granular portfolio construction (Swinkels, 2023; Kreppmeier et al., 2023). Relative to closed-end vehicles, tokenization allows for lower minimum investment thresholds and fewer effective lock-up frictions. Conceptually close to real estate crowdfunding in its single-asset, fractional exposure, tokenization embeds ownership and transfers on interoperable blockchain-based token standards, enabling post-issuance trading and further utilization such as collateralization (Schär, 2021), rather than confining claims within siloed crowdfunding portals with very limited secondary trading options (Lukkarinen and Schvienbacher, 2023). So far, most tokenization initiatives have centered on income-generating investment properties and development project financing, though the infrastructure is also extending to housing-linked financial claims such as home-equity lines of credit (HELOCs), which have already reached substantial scale.² However, active secondary trading remains largely undeveloped, making tokenized investment-property claims the most informative current setting for studying listing behavior and liquidity in blockchain-based real estate markets at scale.

3.2.2 Institutional marketplace details

We study *You and Me (YAM)*, one of the most active and largest peer-to-peer marketplaces for real estate tokens, connected to RealToken LLC (RealT), one of the earliest and largest real estate tokenization platforms by cumulative issuance volume (\$160 million as of 4 March 2026). RealT has tokenized rental residential investment properties in the United States and Latin America since 2019. The tokens are treated as securities under U.S. securities law, consistent with the Howey test, and qualify as investment contracts. Because property deeds generally cannot yet be tokenized directly, each property is held by a dedicated special purpose vehicle (SPV) that retains legal title and operates as a standalone legal entity.³ Property management is predominantly outsourced to local operators contracted by the platform.

Tokens are initially sold through primary issuances comparable to real estate crowdfunding, in which investors complete payment and digitally sign the offering memorandum (see

²For example, Figure’s HELOC tokens have a total asset value of more than \$15 billion.

³Registry-backed pilot projects for direct deed tokenization are beginning to emerge, for example in the UAE, but remain exceptional and early-stage.

Kreppmeier et al. (2023) for details), and are subsequently transferred to investors' digital wallets via blockchain-based smart contracts. To invest in and trade these tokens, users must complete know-your-customer (KYC) verification and be whitelisted by the platform. Token ownership represents in most cases an equity claim on the corresponding SPV, and only more rarely a debt claim. Net rental income, after deduction of platform fees, operating expenses, insurance, and taxes, is distributed weekly to token holders through blockchain transactions settled in stablecoins (i.e., tokens pegged one-to-one to the U.S. dollar), and token holders additionally participate in a property-specific decentralized voting system comparable to an owners' meeting but limited to specified governance matters.

On the secondary market, investors may sell tokens back to the platform at a posted repurchase ("redemption") price tied to the contemporaneous appraisal value and specified balance-sheet items (e.g., maintenance reserves), which provides a continuously visible outside option at a 3% fee comparable to the NAV redemption option of an open-end real estate fund. Alternatively, they can trade directly with one another on peer-to-peer marketplaces.⁴ On the marketplace, sellers post take-it-or-leave-it listings for specified token quantities at stated prices. Listings can be revised or removed at any time, and execution occurs when a buyer accepts the standing terms. Settlement is handled automatically through blockchain-based smart contracts at negligible transaction costs, typically only a few cents, because trades settle on the low-fee Gnosis Chain, an Ethereum-compatible blockchain.⁵ Importantly, price and yield deviations relative to the redemption benchmark are directly visible on the marketplace (see Appendix A.1 for a screenshot). The full sequence of listing actions, including initial posting, price revisions, partial sales, and removals, is publicly traceable at the wallet level through public blockchain records by both market participants and the researcher. These features make the unusually rich transparency in this market a defining characteristic of the setting and a central motivation for our empirical analysis.

Taken together, this institutional setting provides a useful laboratory for studying how real-estate token marketplaces function as tokenization scales. By linking publicly observable listing histories to seller behavior, ToM, and execution outcomes, our findings also speak to the design of conventional online real-estate platforms as they move toward greater transparency for users.⁶ These features shape the underlying trading environment: fractionalization broadens the potential buyer pool, posted execution reduces bargaining frictions, and transparent benchmark pricing makes relative overpricing directly observable. As a result, ToM in tokenized real estate is likely to be more tightly linked to transparent relative pricing, observable market thickness, and publicly traceable seller adjustments than in conventional brokered property markets.

⁴There also exist alternative market designs specific to blockchain technology, aiming to provide instant liquidity (Laschinger et al., 2026). So far, however, these mechanisms have proven only partially effective and highlight the underlying challenge of automated intermediation and liquidity provision in direct real estate markets (Buchak et al., 2025).

⁵Ethereum is a public blockchain network that allows programmable transactions through smart contracts. Gnosis is a compatible blockchain designed for substantially lower transaction fees, but with less market adoption.

⁶For example, major portals such as Zillow increasingly display sale and price-history information as well as listing-history indicators (e.g., "Time on Zillow") on for-sale property pages.

3.3 Related literature and hypothesis development

This section reviews the related literature and develops the hypotheses tested in the empirical analysis.

3.3.1 Related literature

Listing prices, ToM, and selling outcomes in real estate. A large real-estate literature studies ToM as a key measure of market liquidity, emphasizing that sellers face a trade-off between achieving a higher selling price and a faster sale. Early empirical work already treated real-estate sales as a joint price–time problem (Trippi, 1977), and examined whether sellers can trade off sale price against marketing speed through their listing-price strategy (Miller, 1978; Miller and Sklarz, 1987). Theoretical work formalized this trade-off as a selling problem under uncertainty, in which sellers decide whether to accept the currently available offer or continue waiting for a better match, consistent with optimal-stopping logic (Haurin, 1988). Within this framework, the initial listing price is strategically chosen and can serve as an instrument that signals seller valuation and reservation price while shaping bargaining, broker search incentives, and expected ToM (Horowitz, 1992; Yavas and Yang, 1995). Anglin et al. (2003) show that listing-price choice is central to marketing outcomes, as more aggressive initial list prices tend to lengthen ToM by reducing buyer arrival and slowing the matching process. This identifies the listing price, and more specifically the degree of overpricing, as a key mechanism linking ToM to eventual selling outcomes. Other empirical evidence further shows that ToM varies systematically with financing conditions, property characteristics, seller circumstances, and property atypicality (Kang and Gardner, 1989; Glower et al., 1998; Haurin et al., 2010; Cajias and Freudenreich, 2018).

Price revisions, seller adjustment, and listing-history evidence. Beyond the initial list-price decision, theory predicts that sellers adjust prices dynamically when demand is uncertain. In Lazear (1986)’s general price-cutting framework, failure to sell is informative about buyer valuations, implying that sellers may optimally start high and revise downward as they learn from market outcomes. Sass (1988) develops this logic for real estate, showing that optimal price cutting depends on market thinness and the seller’s prior information, while Read (1988) models repricing as a rational learning response to buyer arrival and market feedback. Taylor (1999) further emphasizes that publicly observed market history shapes buyer beliefs, so that ToM and the observable price path become quality signals to which sellers respond strategically. Empirically, Knight (2002) is among the first to incorporate intra-listing revisions into the price–ToM relationship, showing that omitting them biases inference on both sale price and duration. Herrin et al. (2004) extend this by documenting that revision frequency is an economically meaningful margin of seller adjustment and that sellers follow heterogeneous paths, with some making one large cut and others a sequence of smaller ones, though the comparative effectiveness of these strategies remains unresolved. Richer listing-history studies push further: Merlo and Ortalo-Magné (2004) document that revisions are infrequent, typically large, and triggered by prolonged periods without offers; Wit and Klaauw (2013) show that price cuts increase both the selling and withdrawal hazard; and Merlo et al. (2015) document that even small revision costs generate substantial list-price stickiness. Yet all of these studies examine conventional brokered markets where the list price

is only the starting point of negotiation, observed revisions remain an imperfect proxy for true seller adjustment, and the full price-adjustment process is neither equally observable to buyers and sellers nor to the researcher. The transparency provided by blockchain-based systems in tokenized real estate markets remains unmatched and therefore warrants separate attention.

Information transparency, digital intermediation, and tokenized real estate markets. A growing literature shows that information transparency and digital intermediation reshape search, matching, and liquidity in real estate markets. Eerola and Lyytikäinen (2015) show that more detailed public transaction-price information reduces seller uncertainty and improves buyer sorting, raising selling prices and accelerating sales. More recent work finds that online listings and agencies similarly improve search efficiency and transaction performance. Sing and Zou (2024) show that effective online-listing strategies help agents sell faster and above asking price. Gedikli et al. (2026) find that online agencies shorten ToM and lower the sale-list price ratio, and Anderson et al. (2024) report that virtual marketing tools moderate the price–ToM relation. Real estate tokenization can be understood as a stronger form of digital intermediation, changing both the information architecture and the market structure through which real-estate claims trade. Conceptual studies highlight its potential to reduce transaction frictions, enhance transparency, and improve tradability (Smith et al., 2019; Baum, 2021), while also creating legal and governance challenges (Konashevych, 2020; Verstappen, 2025). Empirical work has so far focused on issuance activity, ownership distribution, and fractionalization (Swinkels, 2023; Kreppmeier et al., 2023), portfolio and risk-return implications (Steininger, 2023), secondary market-design challenges (Laschinger et al., 2026), and the potential for tokenization to fill liquidity gaps in traditional real estate markets (Cornelli, 2025). Empirical evidence on listing behavior and marketing time for individual token listings, however, remains largely unexplored.

Taken together, the literature shows that listing prices, ToM, and price revisions are central to selling outcomes, while information transparency and digital intermediation can reshape search, matching, and liquidity. Conventional brokered markets, however, make seller adjustment difficult to observe cleanly because posted prices are blurred by bargaining and revisions are often only partially observed, or not observed at all. Tokenized real-estate marketplaces offer a distinct setting with traceable listing histories, posted-price execution, and a continuously observable valuation benchmark. This allows classic questions on overpricing, seller adjustment, and liquidity to be studied under richer observability, while extending the evidence to tokenized secondary-market listing dynamics. These features directly motivate our hypotheses.

3.3.2 Hypothesis development

Building on the preceding literature, we develop hypotheses for a tokenized real estate marketplace in which listing histories, price revisions, removals, and sales are fully publicly visible, and overpricing is measured relative to the continuously observable redemption price benchmark. We organize the hypotheses in two blocks around two research questions: first, ToM determinants for the full listing sample; and second, seller adjustment and revision outcomes for the subsample of revised listings.

RQ1: What are the determinants of ToM in a tokenized real estate marketplace?

In search-based models of idiosyncratic goods such as real estate, the asking price affects the probability that an arriving buyer accepts the posted terms. A lower asking price expands the set of buyers whose reservation values exceed the ask, thereby increasing the likelihood of a match and shortening expected marketing time (Miller, 1978; Yavas and Yang, 1995; Anglin et al., 2003). In our posted-price marketplace, this mechanism is especially transparent because overpricing can be measured directly relative to an observable valuation benchmark, the redemption price. Listings priced further above this benchmark should therefore remain on the market longer. At the same time, stronger property fundamentals should shorten ToM because they imply more attractive expected cash flows and lower perceived risk, thereby increasing buyer willingness to transact even conditional on the observed premium (Kang and Gardner, 1989; Glower et al., 1998). This motivates hypothesis *H1a*:

H1a (Pricing and ToM)

Listings priced closer to observable valuation benchmarks and backed by stronger property fundamentals are associated with shorter ToM.

Beyond listing-specific characteristics, ToM should also depend on market thickness. In stock-flow matching models, liquidity depends on the interaction between the standing stock of listings and the flow of arriving buyers. Holding demand fixed, a thicker stock of competing listings reduces the likelihood that any given listing matches quickly, thereby increasing expected ToM (Wheaton, 1990; Genesove and Han, 2012). Conversely, stronger buyer flow, proxied by recent trading activity, increases the rate at which investors arrive and screen the standing stock of listings, which should shorten ToM. This leads to hypothesis *H1b*:

H1b (Stock-to-flow and ToM)

A thicker stock of competing listings is associated with longer ToM, while stronger buyer flow is associated with shorter ToM.

Tokenized real estate is also embedded in broader blockchain-based financial infrastructure. Trades and rental distributions are typically settled in stablecoins, and investors can readily reallocate across crypto assets and tokenized real-world assets. Analogous to REITs co-moving with broader stock-market sentiment and liquidity (Ling and Naranjo, 1999; Glascock et al., 2000), tokenized real estate may therefore be exposed to crypto-market conditions through shared investor bases, liquidity, and blockchain-specific frictions. Kreppmeier et al. (2023) show that such market integration is an important determinant of real-estate token issuance success. We therefore expect favorable crypto-market conditions to spill over into secondary trading and shorten ToM. At the same time, because cash flows and valuations are tied to underlying real estate, stronger housing-market conditions should also support demand and shorten ToM (Swinkels, 2023; Steininger, 2023; Bergkamp et al., 2025). By contrast, higher outside yields, such as U.S. Treasury yields, should reduce the relative attractiveness of rent-generating real estate tokens and lengthen ToM. Accordingly, we formulate hypothesis *H1c*:

H1c (Underlying markets and ToM)

More favorable housing- and crypto-market conditions are associated with

shorter ToM for tokenized real estate, whereas higher outside yields are associated with longer ToM.

We next turn to revised listings and ask how publicly observable price-adjustment paths shape marketing times and selling outcomes. In conventional real estate markets, the empirical analysis of seller adjustment is often limited because intermediate list-price revisions and their exact timing are only partially observed, or not observed at all, and because bilateral bargaining can weaken the link between posted and transaction prices (Knight, 2002; Anglin et al., 2003; Merlo and Ortalo-Magné, 2004). In our setting, by contrast, we observe the complete sequence, timing, and magnitude of posted price revisions, and the posted price becomes the transaction price upon acceptance. Thus, the contribution of tokenization here is not transparency in itself, but full traceability, as seller adjustment paths become both empirically observable and potentially informative to buyers (Taylor, 1999). This allows us to study not only whether price revisions accelerate sale, but also whether the path of publicly observable concessions affects final selling outcomes.

This leads to the following research question:

RQ2: How do price revision strategies affect ToM and price outcomes when the full listing path is observable?

We address this question along two related dimensions. First, we examine whether the intensity of a downward price revision affects the remaining time until sale.

This logic extends *H1a* from the initial listing decision to within-listing seller adjustment. In our setting, the relevant mechanism is not only the post-revision price level, but also the publicly traceable revision event itself. A reduction from a 10% premium to 5% should newly bring the listing into the acceptable range for a broader set of previously unmatched buyers than a reduction from 7% to 5%. Larger downward revisions should therefore increase the post-revision matching probability and shorten remaining ToM more strongly than smaller revisions. This motivates hypothesis *H2a*:

H2a (Revision intensity and remaining ToM)

Conditional on the premium remaining after revision, larger downward price revisions reduce remaining ToM.

Second, we examine whether the path by which a given total price concession is delivered affects the premium ultimately achieved at sale. When full price histories are publicly observable, buyers can condition not only on the current listing price but also on the sequence of past revisions. Repeated small downward revisions may signal that a seller's reservation price is gradually declining, thereby encouraging buyers to wait for further concessions (Taylor, 1999). Prior work indicates that sellers follow heterogeneous revision paths, ranging from one relatively large cut to a sequence of smaller cuts (Herrin et al., 2004). However, the comparative effectiveness of these strategies remains largely untested, because such comparisons are difficult to study cleanly in conventional real estate markets. Heterogeneous properties are not directly comparable, revision paths are often only partially observed, and bargaining weakens the link between posted and transaction prices. In our setting, by contrast,

the standardization and fractionalization enabled by real estate tokenization make it possible to compare more directly listings of claims on the same underlying property that arrive at similar total price concessions through different revision paths. If buyers interpret repeated revisions as evidence of a weakening seller, spreading a given concession over many small revisions should reduce the eventual selling premium relative to a single decisive revision. This motivates hypothesis *H2b*:

H2b (Revision path and selling premium)

Holding the total price concession constant, a single decisive price revision achieves a higher final selling premium than multiple smaller revisions.

Importantly, this argument does not require repeated revisions to be irrational. Sellers may revise prices optimally in response to updated beliefs about demand (Lazear, 1986; Read, 1988). Our claim instead concerns buyer inference under transparency. When the full revision path is publicly traceable, successive revisions themselves become information that buyers can use when deciding whether to transact immediately or wait for further concessions, and can therefore be consequential for final selling outcomes. This is the informational channel we test empirically.

3.4 Data

Our empirical analysis combines property-level information on tokenized real estate with complete event histories from blockchain-based secondary-market listings. The sample covers peer-to-peer marketplace activity for fractionalized claims on rental residential investment properties issued by RealT. The traded object in our setting is therefore not the underlying property as a whole, but a transferable claim linked to a property-holding vehicle or its cash flows. Raw marketplace coverage begins on 9 October 2020 with early observations from *Swapcat* and becomes concentrated on *YAM* from 2022 onward, when *YAM* emerges as the dominant peer-to-peer marketplace in the RealT ecosystem. Our final sample ends on 21 July 2025, one day before RealT suspended the redemption mechanism for most properties.⁷ We impose this cutoff because the redemption price is the key valuation benchmark used throughout the paper. Including the subsequent period would mix two distinct market regimes, while the post-suspension window is too short for separate analysis.

We merge three data sources. First, we manually collect property-level information from the RealT website, including location, property characteristics, token supply, and valuation information. Second, we retrieve time-varying property information, in particular rental cash flows and appraisal-related updates, through the publicly available RealT API. Third, we reconstruct complete secondary-market event histories from blockchain data by querying marketplace smart contracts through publicly available blockchain APIs.⁸ These records


⁷The suspension occurred amid broader legal and operational disruptions affecting the platform.

⁸To construct the analysis sample, we exclude private listings, listings with incomplete or inconsistent event histories, and a small number of observations with implausible initial quantities or prices. We also restrict attention to sell-side listings, as buy-side offers are possible on the marketplace but are much less common and not the focus of our analysis. In total, these filters remove roughly 10% of the raw listing data.

contain all events required to reconstruct each listing’s history and price path, together with the corresponding timestamps, prices, quantities, payment tokens, wallet addresses, and token identifiers. This allows us to follow each listing from initial posting until sale, removal, or the end of the sample period, which is essential in our setting, where listings may be revised multiple times, partially sold, or remain active without sale.

We cover 731 tokenized residential real estate assets, predominantly located in the United States and, to a much smaller extent, Latin America. Figure 3.1 maps their geographic distribution. Tokenization activity starts in Michigan, with a pronounced concentration in the Detroit metropolitan area, and subsequently expands to other U.S. markets, first to adjacent Rust Belt locations such as Illinois and Ohio and later to southern states including Texas and Florida. RealT eventually also extends tokenization to Latin America, including Mexico, Panama, and Colombia, although activity there remains limited as of July 2025. This spatial rollout is consistent with broader evidence that digital innovations often emerge in concentrated hubs before diffusing more gradually across regions (Gedikli et al., 2026; Kalyani et al., 2025).

Figure 3.1: Geographic distribution of RealT’s tokenized real estate across the U.S. and Latin America



Note: Latin America is shown at the country level (Colombia, Mexico, Panama); the U.S. at the state level (Alabama, Florida, Georgia, Illinois, Indiana, Michigan, Missouri, New York, Ohio, Texas). Data as of July 2025.

Panel A of Table 3.1 summarizes the underlying property and locational characteristics. The average tokenized asset has an initial investment value of roughly \$200,000, annual net rent of about \$16,000, and a token supply of around 4,000 units, although these distributions are strongly right-skewed, with the largest tokenized multi-family property valued at roughly \$1.9 million. Most assets are small-scale residential properties. Around three quarters are single-family homes, the average number of rental units is three, and the median property has only one unit. The sample is also relatively old, with an average building age of 78 years, and around 10% of properties are classified as renovation required, which we define based

on publicly disclosed maintenance amounts at issuance of at least 15% of total investment value. As indicated in Figure 3.1, roughly three quarters of all tokenized assets in the sample are located in Detroit. The sample is also relatively urban, with an average distance to the central business district of about seven miles. In terms of token structure, a small share of observations (10%) consists of portfolio tokens, which resemble pooled vehicles rather than single-asset tokenizations, and the overwhelming majority represent equity claims on the underlying property (98.5%).

Panel B of Table 3.1 reports time-varying market conditions. Here, the number of observations corresponds to the number of sample days. To capture these conditions, we merge the token-level data with external series reflecting local housing, crypto, and yield environments, as described in the table notes.

We organize the marketplace data into two complementary empirical units. The first is the listing level, where we observe each posted listing from its creation until sale, removal, or the end of the sample period. This level is used to study the determinants of overall ToM and the relationship between listing strategies and selling outcomes, corresponding to *H1* and *H2b*. The second is the spell level, where we decompose a listing into intervals that begin with a price revision and end at the next sale, removal, further revision, or censoring point. This within-listing perspective allows us to examine how a specific revision affects the remaining ToM, corresponding to *H2a*.⁹

Figure 3.2 illustrates this distinction. The step line traces the listing premium relative to the contemporaneous redemption price over the lifetime of one exemplary listing. At the listing level, we measure ToM from initial posting to the first sale, which is the closest analogue to the standard notion of marketing time in the real-estate literature and provides a tractable measure in a setting with potentially partial executions. This approximation is empirically meaningful because the median first-sale share in our sample is 100%, while the mean is about 66% (see Table 3.2).¹⁰ At the spell level, we restart the clock at each downward revision and measure the remaining time until the subsequent sale. This distinction is central for separating the effect of the initial pricing decision from the effect of later seller adjustments.

Table 3.2 summarizes the main empirical variables and subsamples used throughout the analysis. Panel A focuses on the listing level and distinguishes between all listings, sold listings, removed listings, revised listings, and sold revised listings. This mirrors the structure of the listing-level analyses: the full listing sample underlies the ToM specifications in *H1*, while the subsample of sold revised listings forms the basis for the revision-strategy analysis in *H2b*. In addition to the core regression variables, such as ToM, listing premium relative to redemption price, listed volume, competing listings volume, prior-day token sales, and time since appraisal, the table also reports contextual variables that help interpret the divisible nature of listings, in particular selling price, redemption price, first-sale share, and total price concessions. Panel B turns to the spell level and summarizes all downward revision spells as well as the subset ending in sale. This is the relevant empirical unit for *H2a*.

⁹We also observe upward price revisions, which account for 3,330 of 26,518 price revisions (12.6%). Because they are comparatively rare and economically distinct from downward repricing, we exclude them from the spell-level revision analysis and focus on downward revisions throughout.

¹⁰As a robustness check, we also redefine listing-level ToM using an inventory-clearance threshold. Appendix A.2 reports the corresponding estimates for time until 90% of the initially listed quantity is sold.

Table 3.1: Summary statistics for property and market characteristics

	N	Mean	SD	Min	Median	Max
Panel A: Property and locational characteristics						
Property investment value USD	731	201,202.78	270,850.38	48,080.00	83,671.50	1,881,450.00
Annual net rent USD	731	16,076.84	23,698.22	0.00	6,909.69	185,510.06
Total tokens	731	3,929	5,286	750	1,650	37,000
Annual net rent per token USD	731	4.24	1.73	0.00	4.42	19.29
Distance to CBD (miles)	731	7.36	3.45	0.00	7.34	44.17
Building age (years)	731	78	29	0	84	145
Interior size (sqft)	731	2,986	4,500	667	1,335	45,305
Rental units	731	3	6	1	1	58
Single family	731	75.65%				
Multi family	731	14.77%				
Duplex	731	9.58%				
Portfolio token	731	9.17%				
Real Estate Loan	731	1.50%				
Real Estate Equity	731	98.50%				
Development project	731	1.50%				
Renovation required	731	10.40%				
Detroit	731	75.92%				
Cleveland	731	8.48%				
Chicago	731	2.74%				
St. Louis	731	1.92%				
Toledo	731	1.23%				
Garfield Heights	731	1.09%				
Highland Park	731	0.82%				
Montgomery	731	0.68%				
Maple Heights	731	0.41%				
Euclid	731	0.41%				
Other cities	731	6.29%				
Panel B: Market conditions						
Housing market	1,747	0.006	0.005	-0.006	0.005	0.021
Crypto market	1,747	0.004	0.067	-0.318	0.003	0.426
Crypto market shock	1,747	0.16	0.37	0.00	0.00	1.00
1-month U.S. Treasury yield (%)	1,747	3.41	2.19	0.00	4.37	6.02


Note: This table reports summary statistics for property and market characteristics. Panel A covers the 731 tokenized assets in the sample. Annual net rent is measured as the time-weighted average from token issuance to the sample end, and annual net rent per token is obtained by dividing this measure by total token supply. Panel B reports daily market conditions over the sample period (9 October 2020 to 21 July 2025), so the number of observations corresponds to the number of sample days. The housing market variable is the equally weighted monthly change in the relevant house-price index across sample locations, using S&P Case-Shiller Home Price Indices for the corresponding U.S. metropolitan areas and housing-related consumer price index components for the Latin American countries, as comparable monthly real estate price indices are not consistently available. Crypto market refers to the prior 3-day cumulative Ether return, crypto market shock indicates a negative prior 3-day Ether return of at least 5%, and the treasury variable is the daily 1-month U.S. Treasury yield. Dummy variables are reported as percentages.

Table 3.2: Summary statistics for main regression variables

	N	Mean	SD	P10	Median	P90
Panel A: Listing-level						
<i>All listings</i>						
ToM (sold/removed/end of sample)	111,180	11.33	34.53	0.00	1.10	27.76
Listing premium (%) relative to redemption price	111,180	12.48	19.35	-0.12	6.88	32.59
Listing price USD	111,180	59.32	17.74	50.15	54.88	70.75
Redemption price USD	111,180	52.64	12.03	50.00	50.64	54.45
Listed volume USD	111,180	559.56	4,430.11	23.59	101.15	822.01
Competing listings volume USD	111,180	4,441.11	14,534.08	0.00	629.25	9,428.22
Prior-day token sales USD	111,180	259.84	1,800.65	0.00	14.10	478.03
Time since appraisal (days)	111,180	491.99	382.55	39.00	420.00	1,076.00
<i>Sold listings</i>						
First-sale share	72,917	0.66	0.41	0.02	1.00	1.00
ToM (sold)	72,917	6.31	25.82	0.00	0.31	13.65
Listing premium (%) relative to redemption price	72,917	9.22	13.98	-0.47	4.81	25.25
Selling price USD	72,917	57.20	15.77	50.00	53.00	66.31
<i>Removed listings</i>						
ToM (removed)	35,007	19.11	44.25	0.12	4.00	49.91
Listing premium (%) relative to redemption price	35,007	19.33	26.42	0.00	10.09	53.00
<i>Revised listings</i>						
Listing premium (%) relative to redemption price	11,961	12.72	16.07	0.00	7.96	29.73
<i>Sold listings (after price revision)</i>						
Selling premium (%) relative to redemption price	9,612	8.16	13.28	-1.65	3.99	23.22
Single price-revision strategy	9,612	0.59	0.49	0.00	1.00	1.00
Price revisions: 2	9,612	0.20	0.40	0.00	0.00	1.00
Price revisions: 3-4	9,612	0.13	0.34	0.00	0.00	1.00
Price revisions: 5+	9,612	0.08	0.26	0.00	0.00	0.00
Total price concession (%)	9,612	5.03	6.31	0.59	3.13	11.22
Panel B: Spell-level						
<i>All price revisions</i>						
Remaining ToM (sold/removed/end of sample)	23,188	5.07	21.00	0.00	0.80	10.08
Price revision (%)	23,188	2.40	4.27	0.12	1.46	5.21
Listing premium at revision (%) relative to redemption price	23,188	9.02	13.90	-1.25	5.01	23.95
Number of revisions	23,188	7.48	16.71	1.00	3.00	15.00
<i>Sold revisions (sold spells)</i>						
Remaining ToM (sold)	10,880	3.32	13.68	0.00	0.31	6.95
Price revision (%)	10,880	2.77	4.10	0.29	1.76	6.20
Listing premium at revision (%) relative to redemption price	10,880	6.88	12.11	-2.03	3.20	20.53
Number of revisions	10,880	3.40	8.42	1.00	2.00	6.00

Note: This table reports summary statistics for the main regression variables. Panel A covers listing-level variables, including all listings, sold listings, removed listings, revised listings, and sold listings after price revision. Panel B reports spell-level variables for all price revisions and sold revisions. All price revisions are downward price revisions. Price-revision strategy variables are indicator variables, and first-sale share measures the share of listed tokens sold in the first purchase.

Figure 3.2. Listing price path: listing-level and spell-level (illustrative example)



Note: The step line shows the listing premium relative to the redemption price over ToM (days) for one exemplary listing (ID 34129; REALTOKEN-S-9576-GRANDMONT-AVE-DETROIT-MI). Triangles mark downward list-price revisions; labels report the revision magnitude in percentage points. Stars indicate sales. The vertical dashed line marks the first sale after 15.4 days. The bold segments “Spell 1” and “Spell 2” illustrate the spell-level construction, where remaining ToM is measured from a downward price revision to the subsequent sale; in this example, remaining ToM equals 5.7 days for Spell 1 and 4.5 days for Spell 2. For illustration, we display only the first two purchases even though this listing continues afterwards, as indicated by the short dashed segment after the second sale.

Key variables are constructed as follows. Listing-level ToM is measured as the number of days from a listing’s initial posting to its first sale, removal, or the end of the sample period, depending on the relevant subsample. For revision spells, remaining ToM is measured from a downward price revision to the subsequent sale, removal, further revision, or censoring point. Listing premium relative to redemption price is computed as the percentage difference between the posted listing price and the contemporaneous redemption price, which provides a continuously observable valuation benchmark in our setting. Competing listings volume (USD) measures contemporaneous supply-side depth for the same underlying property, defined as the USD value of all other active listings at the listing’s creation timestamp, excluding the focal listing. Prior-day token sales (USD) capture lagged demand by measuring the USD trading volume of the same underlying token on the previous day.


Several descriptive patterns stand out. First, listing durations are short in absolute terms: across all listings, median ToM is 1.10 days, while the median ToM among sold listings is only 0.31 days.¹¹ These magnitudes should not be interpreted as marketing times for whole properties, but for small, fractionalized real-estate claims traded at negligible transaction costs. With a median listed volume of only about \$101 (roughly two tokens at the median list price of \$54.88), a typical listing represents just about 0.12% of the median underlying property investment value of \$83,671, which helps explain why matching can occur much faster than in conventional property markets by making participation feasible for a broader set of buyers and by reducing the economic stakes of any individual purchase. By contrast,

¹¹This pattern is robust to a stricter sale threshold requiring 90% of the initially listed quantity to be sold, as applied in Appendix A.2, yielding a median ToM of 0.47 days and a mean ToM of 10.07 days.

unsuccessful listings remain substantially longer on the market and are posted at materially higher premia relative to redemption price. Removed listings have a median ToM of 4.00 days and a median listing premium of 10.09%, compared with a median listing premium of 4.81% for sold listings. This descriptive contrast already suggests that relative pricing is closely linked to execution outcomes.

Overall, price revisions are relatively infrequent but economically meaningful. Only 11,961 out of 111,180 listings are revised at least once, indicating substantial listing-price stickiness despite negligible mechanical adjustment costs as theoretically reasoned by Merlo et al. (2015). Among sold revised listings, the average total price concession is about 5% relative to the initial list price, and the average selling premium over redemption price remains positive at 8.16%. Figure 3.3 further shows that most sold revised listings follow relatively simple repricing paths and generally modest total cumulative price concessions, with a majority involving only a single downward revision and progressively fewer listings exhibiting repeated cuts. This heterogeneity in revision paths and overall concessions motivates the distinction between our spell-level analysis of remaining ToM after a given revision (*H2a*) and our listing-level analysis of revision strategies and eventual selling premia (*H2b*).

Figure 3.3. Distribution of revision frequency and total price concessions per sold listing



Note: The sample is restricted to sold listings that experienced at least one downward price revision. Panel (a) shows the distribution of the number of downward price revisions per sold revised listing. Panel (b) shows the distribution of the total cumulative price concession, measured relative to the initial listing price. The number of observations is $N = 9,612$ in both panels.

Within listings, revisions are also followed by comparatively short remaining durations. Across all revision spells, median remaining ToM is 0.80 days, and among sold revision spells it is only 0.31 days. At the same time, sold spells are associated with somewhat larger average revision sizes and lower remaining premia than the full set of revision spells. This pattern is consistent with the idea that downward repricing helps move listings toward execution, while still leaving meaningful cross-sectional variation in revision intensity and post-revision liquidity.

3.5 Methodology

Our empirical analysis documents systematic relationships between listing strategies, ToM, and selling outcomes in a tokenized real estate marketplace with fully traceable listing

histories.

The methodology follows the structure of our hypotheses and combines two complementary approaches. First, we estimate OLS models on sold observations, which provide a simple and transparent benchmark for how listing characteristics and price-setting choices relate to observed ToM and selling premia. Second, we estimate duration models that incorporate censored observations and therefore exploit a broader part of the listing history. This is important in our setting because restricting attention to sold listings alone may yield a selected view of market liquidity by discarding listings that are removed or remain unsold at the end of the sample. Because ToM is a time-to-event outcome and some observations are censored, duration models provide the natural complementary framework and are commonly used in related real-estate ToM research, often in hazard form (see, e.g., Anglin et al., 2003; Wit and Klaauw, 2013).

Although Cox models are among the most common variants in related ToM studies, formal diagnostics reject the proportional-hazards assumption in our data, indicating that the effects of key covariates are not constant over the marketing period. We therefore use accelerated failure time (AFT) models as our main duration framework. In such cases, AFT models provide a useful alternative (Wei, 1992; Kalbfleisch and Prentice, 2002) and have also been applied in related real-estate and tokenized-asset settings (Carrillo and Williams, 2019; Kreppmeier et al., 2023). This choice is also attractive substantively, since AFT coefficients map directly into the question of whether a covariate lengthens or shortens expected ToM. We estimate both lognormal and loglogistic specifications, which are well suited to the strongly right-skewed duration outcomes in our data and provide greater flexibility than the Weibull alternative commonly used in related studies.

In the AFT framework, the logarithm of duration is modeled as

$$\log(T_i) = L_i' \beta_L + P_i' \beta_P + M_t' \beta_M + \alpha_q + \alpha_c + \varepsilon_i, \quad (3.1)$$

where T_i denotes the relevant duration outcome for observation i , L_i is a vector of listing- or spell-level characteristics, P_i is a vector of property characteristics, and M_t is a vector of market conditions. The terms α_q and α_c denote quarter-year and city fixed effects, respectively, while ε_i follows a lognormal or loglogistic distribution depending on the specification. We report $\exp(\beta)$ as time ratios, so values above one indicate longer expected duration and values below one indicate shorter expected duration.

The empirical design has three building blocks. For *H1*, we study listing-level ToM from listing creation to first sale using both sold-only OLS specifications and AFT models that treat removals and end-of-sample observations as right-censored. For *H2a*, we move to the spell level and study the remaining ToM after a downward price revision, again combining sold-only OLS with AFT duration models that treat non-sale spell endings as censored. For *H2b*, we return to the listing level and use OLS to examine whether different revision paths are associated with different final selling premia for sold listings with at least one downward price revision. Across specifications, we include property and market controls together with time fixed effects. The AFT models include city fixed effects throughout, while the sold-only OLS specifications alternatively use city or underlying-property fixed effects. Standard errors are clustered at the underlying-property level.

A further methodological difference from much of the traditional real estate literature concerns

the measurement of overpricing. In many conventional real estate studies, overpricing is approximated indirectly by first estimating a typical or normal listing price and then interpreting the residual from that benchmark as over- or underpricing (see, e.g., Anglin et al., 2003; Wit and Klaauw, 2013; Merlo and Ortalo-Magné, 2004). In our setting, we do not need to rely on such proxies. Instead, we measure listing and selling premia directly relative to the redemption price at which the platform stands ready to repurchase the token. This benchmark is economically meaningful because it represents a visible outside option for market participants and provides a common reference point for both sellers and buyers when evaluating secondary-market prices.

3.6 Results

This section presents the empirical evidence for the hypotheses developed in Section 3.3.2.

3.6.1 Determinants of ToM in a tokenized real estate marketplace

Table 3.3 reports the main evidence on the determinants of ToM for tokenized real estate listings. Columns (1) and (2) present sold-only OLS estimates, while Columns (3) and (4) report AFT models that treat sales as the event of interest and removals as censored observations.¹² We focus on the AFT estimates because they incorporate incomplete listing spells and provide directly interpretable time ratios, where values above one indicate longer ToM and values below one shorter ToM. The main findings are robust to redefining ToM as the time until 90% of the initially listed quantity is sold rather than the time to first sale (Appendix Table A.2).

H1a: Pricing and ToM. Consistent with *H1a*, listing overpricing relative to the redemption price is one of the strongest and most robust predictors of ToM. In the AFT models, a one-percentage-point increase in the listing premium is associated with time ratios of 1.17 and 1.19, implying roughly 17–19% longer expected ToM. Evaluated at the sample mean ToM of about 11 days, this corresponds to roughly one and a half additional days on market. This aligns with search-based theories of real estate markets, in which more overpriced listings reduce buyer arrival rates and prolong marketing time (Yavas and Yang, 1995; Anglin et al., 2003; Haurin et al., 2010). Compared with conventional real estate markets, where overpricing is typically inferred from noisier hedonic benchmarks, tokenization makes relative overpricing more salient for both market participants and the researcher.

Annual net rent per token is a stable and economically meaningful predictor of faster sale. The AFT time ratios of 0.82 and 0.81 imply that a \$1 increase in annual net rent per token, equivalent to roughly one quarter of the mean annual rent of about \$4 per token, is associated with about 18–19% shorter expected ToM. Physical characteristics add little incremental explanatory power once pricing is controlled for. Development status, renovation need, building age, and downtown distance are mostly weak or inconsistent, with real estate loans

¹²Appendix Table A.1 complements the main sales analysis by examining ToM until removal. The results show that removals are far less sensitive to overpricing and more strongly shaped by broader market conditions, consistent with removal representing a distinct seller adjustment decision (Wit and Klaauw, 2013; Taylor, 1999).

as the main exception, suggesting that these instruments are priced and traded differently. This pattern is consistent with the market-efficiency intuition in Yavas and Yang (1995): once posted prices already incorporate much of the information contained in observable property characteristics, those characteristics add little additional explanatory power for ToM. This mechanism may be even stronger because key property information is standardized, digitally displayed, and commonly observable to both buyers and sellers, reducing bilateral information frictions related to visible asset quality. Cash-flow attractiveness, by contrast, remains directly relevant for investors facing heterogeneous opportunity costs of capital. This also suggests that ToM in tokenized real estate may be less informative about hidden physical quality and more informative about pricing and market absorption (Taylor, 1999).

H1b: Stock-to-flow and ToM. Table 3.3 also strongly supports *H1b*. Thicker supply-side depth is associated with slower matching, while stronger buyer flow is associated with faster sales. In the AFT models, higher competing listings volume is associated with longer expected ToM, whereas stronger prior-day token sales are associated with substantially shorter expected duration. Higher own listed volume is likewise associated with longer ToM, suggesting that larger sell orders take longer to absorb. A doubling of competing listings volume is associated with roughly 12% longer expected ToM, while a doubling of prior-day token sales is associated with roughly 37–38% shorter expected ToM. This pattern is consistent with stock–flow matching intuition. Holding demand fixed, a larger stock of competing offers lowers the sale probability of any individual listing, whereas stronger recent trading signals active buyer demand (Wheaton, 1990; Genesove and Han, 2012). Rather than bilateral offers for an entire property, liquidity is formed in a digital marketplace where competing supply and realized transaction flow for the same token can be observed at high frequency and low denomination. Tokenization therefore does not eliminate real-estate search frictions, but shifts them into a more transparent microstructure in which relative pricing and platform-level stock and flow become directly measurable.

H1c: Underlying markets and ToM. The broader market variables present a mixed but informative picture for *H1c*. Stronger housing-market conditions are associated with shorter ToM, indicating that conditions in the underlying real-estate market continue to matter (Swinkels, 2023; Bergkamp et al., 2025). This is also consistent with Carrillo (2013), who shows that selling outcomes jointly reflect underlying housing-market conditions. The 1-month U.S. Treasury yield enters positively in most specifications, consistent with higher outside yields reducing the relative attractiveness of rent-generating real-estate tokens. This parallels the broader real-estate literature, in which financing conditions and interest rates affect marketing time (Ferreira and Sirmans, 1989; Kang and Gardner, 1989). The crypto-market variables, however, do not support the ex ante prediction that more favorable crypto conditions shorten ToM. Instead, stronger crypto-market conditions are associated with longer durations, which could be interpreted as reflecting investor attention and capital shifting toward higher-volatility crypto opportunities. At the same time, the positive and highly significant crypto-shock effect indicates that adverse crypto-market disruptions also slow trading (Kreppmeier et al., 2023). These results suggest that ToM of tokenized real estate is shaped by both its underlying property market and the broader blockchain-based investment environment.

Overall, the evidence supports *H1* primarily through the overpricing and stock–flow channels. Higher premia and thicker competing inventory are associated with longer ToM, whereas stronger same-token trading activity and higher rental cash flows are associated with faster

matching. More broadly, the results suggest that tokenization does not remove the slow matching typical of real estate, but changes how it is expressed: ToM becomes more tightly linked to transparent relative pricing, observable platform-level liquidity, and the broader opportunity set of blockchain-based investors.

Table 3.3: Determinants of Time on Market (ToM) – Sales

	Dependent variable:			
	ToM until sale (sold only) <i>OLS</i>		ToM until sample end (removals censored) <i>AFT</i>	
	(1)	(2)	<i>lognormal</i> (3)	<i>loglogistic</i> (4)
Constant	0.43 (0.59)	1.47*** (0.54)	3.18*** (1.06)	3.78*** (1.22)
Listing premium (%) relative to redemption price	0.12*** (0.01)	0.16*** (0.01)	1.17*** (0.00)	1.19*** (0.00)
ln(1+ Listed volume USD)	0.23*** (0.02)	0.21*** (0.02)	1.18*** (0.01)	1.17*** (0.01)
Time since appraisal (days)	-0.00*** (0.00)	-0.00*** (0.00)	1.00*** (0.00)	1.00*** (0.00)
ln(1+ Competing listings volume USD)	0.10*** (0.01)	0.10*** (0.01)	1.18*** (0.01)	1.18*** (0.01)
ln(1+ Prior-day token sales USD)	-0.60*** (0.01)	-0.55*** (0.01)	0.50*** (0.00)	0.51*** (0.00)
Annual net rent per token (USD)	-0.17*** (0.06)	-0.23*** (0.02)	0.82*** (0.01)	0.81*** (0.01)
Rental units	0.03*** (0.01)		1.04*** (0.00)	1.04*** (0.00)
Development project	-0.85 (0.68)		1.03 (0.33)	1.25 (0.40)
Renovation required	0.02 (0.15)		1.10 (0.06)	1.05 (0.06)
Real estate loan	-1.00*** (0.38)		0.29*** (0.06)	0.29*** (0.07)
Building age (years)	-0.00 (0.00)		1.00* (0.00)	1.00 (0.00)
Downtown distance (miles)	-0.00 (0.02)		1.01** (0.01)	1.01 (0.01)
Housing market	-0.18*** (0.05)	-0.18*** (0.05)	0.66*** (0.03)	0.63*** (0.03)
Crypto market	0.75*** (0.26)	0.81*** (0.25)	1.64* (0.43)	1.61* (0.41)
Crypto market shock	0.19*** (0.05)	0.27*** (0.05)	1.29*** (0.06)	1.29*** (0.06)
1-month US Treasury yield	0.18** (0.08)	0.26*** (0.08)	1.13* (0.08)	1.09 (0.08)
Quarter-year FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	Yes
Property FE	No	Yes	No	No
Observations	72,917	72,917	111,180	111,180
R^2	0.33	0.39		
Adjusted R^2	0.33	0.38		
Log-likelihood			-78072.5	-77847.7
χ^2 (df = 79)			39021.9	41169.7

Notes: The outcome variable is ToM (T), measured as the time from listing creation until the first sale. Because listings may be partially executed, first sale provides the relevant exit point for the main analysis. All models are estimated on observations with $T > 0$. Columns (1)–(2) report OLS estimates for sold listings, with dependent variable $\log(T)$ and cluster-robust standard errors at the underlying-property level. Column (1) includes quarter-year and city fixed effects. Column (2) includes quarter-year and underlying-property fixed effects. Columns (3)–(4) report AFT models of time to first sale. Removed listings and listings still active at the end of the sample are treated as right-censored. In the AFT sample, the $T > 0$ restriction excludes 548 observations, all of which correspond to sales at $T = 0$. AFT columns report time ratios, $\exp(\beta)$, with delta-method standard errors in parentheses, where $\text{se}(\exp(\beta)) \approx \exp(\beta) \text{se}(\beta)$. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

3.6.2 Price revision strategies under full observability

We now turn to revised listings and examine how publicly observable price-adjustment paths shape subsequent matching and final sale outcomes. We first analyze whether larger downward revisions shorten remaining ToM after a revision (*H2a*). We then examine whether, holding the total price concession constant, the path by which that concession is delivered affects the final selling premium (*H2b*).

H2a: Revision intensity and remaining ToM. Table 3.4 provides evidence that larger downward price revisions shorten the remaining time until sale, conditional on how overpriced the listing still is after the revision. In both AFT specifications, a one-percentage-point larger price revision is associated with a time ratio of 0.94, implying about 6% shorter expected remaining ToM. The effect is precisely estimated and stable across specifications. Evaluated at the average revision size of 2.4%, this corresponds to roughly 14% shorter remaining duration. Importantly, this result holds after controlling for the listing premium that remains at the time of revision. The premium at revision itself enters with time ratios of 1.14 and 1.16, confirming that listings that remain substantially overpriced even after a revision still take longer to sell. Taken together, these results support the search-based prediction that a larger downward price revision provides a stronger demand stimulus, while also showing that buyers continue to respond to the degree of residual overpricing, consistent with earlier evidence that list-price reductions increase the selling rate and that their timing and magnitude matter (Wit and Klaauw, 2013). This joint response to revision intensity and residual overpricing is precisely the dimension that is difficult to isolate in conventional real estate settings (Knight, 2002).

The number of revisions is also informative, but in a different way. In the AFT models, each additional revision is associated with time ratios of 1.05 and 1.06, implying about 5–6% longer expected remaining duration. Rather than suggesting that repeated revisions accelerate sale, this pattern is more consistent with repeated revisions characterizing harder-to-clear listings with persistently weak market absorption. Put differently, larger price revisions are associated with shorter remaining ToM, but longer revision histories are associated with more persistent illiquidity. The evidence therefore supports *H2a* in its core revision-intensity dimension, while indicating that repeated revisions are a marker of difficult-to-sell listings rather than a clean acceleration mechanism.

Several control variables suggest that post-revision matching remains shaped by the same broad liquidity forces identified in *H1*. Remaining ToM after a revision is longer when competing supply is higher and shorter when recent token demand and cash-flow attractiveness are stronger. Housing-market conditions and short-term interest rates also continue to matter, whereas the broader crypto market is no longer robust in the AFT specifications.

Table 3.4: Determinants of remaining ToM after price revision


	Dependent variable: Remaining ToM after price revision			
	ToM until sale (sold only) <i>OLS</i>		ToM until sample end (removals/revisions censored) <i>AFT</i>	
	(1)	(2)	<i>lognormal</i> (3)	<i>loglogistic</i> (4)
Constant	-3.07*** (1.11)	-3.70*** (1.15)	6.24*** (3.70)	5.29*** (3.08)
Price revision (%)	-0.08*** (0.01)	-0.06*** (0.02)	0.94*** (0.01)	0.94*** (0.01)
Listing prem at revision (%) rel to redemption price	0.07*** (0.01)	0.12*** (0.01)	1.14*** (0.00)	1.16*** (0.00)
Number of revisions	-0.01 (0.01)	-0.01 (0.01)	1.05*** (0.01)	1.06*** (0.01)
ln(1+ Listed volume USD)	0.09*** (0.02)	0.06*** (0.02)	0.94*** (0.02)	0.91*** (0.02)
Time since appraisal (days)	-0.00*** (0.00)	-0.00** (0.00)	1.00*** (0.00)	1.00*** (0.00)
ln(1+ Competing listings volume USD)	0.03*** (0.01)	0.03** (0.01)	1.07*** (0.01)	1.07*** (0.01)
ln(1+ Prior-day token sales USD)	-0.59*** (0.01)	-0.56*** (0.02)	0.48*** (0.01)	0.51*** (0.01)
Annual net rent per token (USD)	-0.10*** (0.03)	-0.23*** (0.03)	0.87*** (0.01)	0.86*** (0.01)
Rental units	0.03*** (0.01)		1.05*** (0.00)	1.05*** (0.00)
Development project	-1.08*** (0.34)		0.29* (0.19)	0.22** (0.14)
Renovation required	0.19 (0.13)		1.13 (0.14)	1.07 (0.13)
Real estate loan	-0.27 (0.40)		1.30 (0.50)	1.24 (0.48)
Building age (years)	-0.00* (0.00)		1.00*** (0.00)	1.00** (0.00)
Downtown distance (miles)	-0.01 (0.01)		0.96*** (0.01)	0.95*** (0.01)
Housing market	-0.29*** (0.11)	-0.39*** (0.12)	0.72*** (0.08)	0.77** (0.08)
Crypto market	1.44*** (0.48)	1.18** (0.51)	1.40 (0.73)	1.04 (0.51)
Crypto market shock	-0.00 (0.11)	0.09 (0.11)	0.97 (0.11)	0.94 (0.10)
1-month US Treasury yield	0.36** (0.16)	0.53*** (0.17)	1.54*** (0.19)	1.57*** (0.20)
Quarter-year FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	Yes
Property FE	No	Yes	No	No
Observations	10,880	10,880	23,188	23,188
R^2	0.29	0.36		
Adjusted R^2	0.28	0.31		
Log-likelihood			-14377.0	-14257.4
χ^2 (df = 73)			6372.5	6479.9

Notes: The outcome variable is remaining ToM (T), measured as the time from a downward price revision until the next sale. All models are estimated on revision spells with $T > 0$. Columns (1)–(2) report OLS estimates for sold revision spells, with dependent variable $\log(T)$ and cluster-robust standard errors at the underlying-property level. Column (1) includes quarter-year and city fixed effects. Column (2) includes quarter-year and underlying-property fixed effects. Columns (3)–(4) report AFT models of time to sale after a downward price revision. Revision spells ending in a further revision, removal, or the end of the sample are treated as right-censored. AFT columns report time ratios, $\exp(\beta)$, with delta-method standard errors in parentheses, where $\text{se}(\exp(\beta)) \approx \exp(\beta) \text{se}(\beta)$. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

H2b: Revision path and selling premium. Table 3.5 examines whether the path by which a given total price concession is delivered affects the final selling premium over redemption price. The central result is that, holding the total concession fixed, gradual multi-step repricing is associated with worse execution outcomes. In Columns (1)–(2), the coefficient on the single-revision strategy indicator is positive and highly significant throughout. Relative to listings sold after multiple price revisions, listings sold after exactly one downward price revision achieve final selling premia that are about 1.5–1.8% higher. This supports *H2b* and indicates that concentrating the concession in one decisive price revision is associated with a better selling outcome than spreading the same adjustment over multiple revisions.

Columns (3)–(4) show that this effect is monotonic in the number of revisions. Relative to listings sold after exactly one downward price revision, listings sold after two revisions achieve premia that are about 1.1–1.3% lower, listings with three to four revisions about 1.6–1.9% lower, listings with five to six revisions about 2.3–2.6% lower, listings with seven to eleven revisions about 3.7–4.3% lower, and listings with twelve or more revisions about 4.9–5.6% lower. The negative gradient is economically meaningful and highly systematic, implying that execution premia decline as the public revision history becomes more protracted. Figure 3.4 visualizes this monotonic gradient for the Column (4) specification, showing that selling premia decline steadily as the number of price revisions increases.

Figure 3.4. Effect of price revision path on final selling premium



Note: This figure plots the estimated effect of the publicly traceable price revision path on the final selling premium over redemption price, relative to listings sold after exactly one downward price revision. Coefficients are based on Column (4) of Table 3.5. Holding the total price concession constant, final selling premia decline monotonically as the number of price revisions increases. Vertical bars indicate 95% confidence intervals.

Importantly, these specifications compare listings with the same overall concession. By controlling for total price concession and the premium at the first revision, the estimates isolate the informational content of the revision path rather than simply the level of price adjustment.

The evidence is consistent with buyers conditioning on the publicly traceable revision history rather than only on the final ask. Repeated small price revisions may reveal a seller whose

reservation price is progressively declining, thereby strengthening buyer incentives to wait for further concessions. The full revision path has independent informational content, consistent with the broader idea that publicly observable market history shapes buyer inference and seller outcomes (Taylor, 1999). Relative to conventional housing markets, this mechanism is especially sharp in our setting, as the posted price is the execution price upon acceptance, rather than the starting point for subsequent bargaining (Merlo and Ortalo-Magné, 2004). Therefore, buyers respond not only to how much sellers concede (*H2a*), but also to how those concessions are delivered over time (*H2b*).

Among the controls, total price concession is negatively associated with the final selling premium, as expected, while the premium at the first revision is positively associated with the final outcome, indicating that listings starting from a higher premium at the first revision also tend to retain more premium at sale. Broader housing- and crypto-market variables are mostly insignificant in these specifications, suggesting that within the subsample of sold revised listings, execution premia are driven primarily by seller-specific pricing paths rather than by aggregate market conditions. The high R^2 values are not surprising, since the dependent variable is itself a final premium measure. The central result is that the revision-path coefficients remain economically large, monotonic, and statistically significant even after conditioning on the price-level components. This pattern is also robust to using the selling price, rather than the selling premium over redemption price, as the dependent variable. Appendix Table A.3 shows the same qualitative results. Overall, Table 3.5 provides strong support for *H2b*, indicating that when the full listing path is observable, a single decisive price revision is associated with a higher ultimate selling premium compared to gradual multi-step price reductions.

Table 3.5: Revision strategy and selling premium

	Dependent variable: Selling premium relative to redemption price			
	(1)	(2)	(3)	(4)
Single-revision strategy (=1)	0.018*** (0.001)	0.015*** (0.001)		
Price revisions: 2 (vs 1)			-0.013*** (0.001)	-0.011*** (0.001)
Price revisions: 3-4 (vs 1)			-0.019*** (0.002)	-0.016*** (0.001)
Price revisions: 5-6 (vs 1)			-0.026*** (0.003)	-0.023*** (0.003)
Price revisions: 7-11 (vs 1)			-0.043*** (0.005)	-0.037*** (0.005)
Price revisions: 12+ (vs 1)			-0.056*** (0.009)	-0.049*** (0.008)
Total price concession (%)	-0.001*** (0.000)	-0.001** (0.000)	-0.001* (0.000)	-0.001* (0.000)
Premium at first revision (%)	0.007*** (0.000)	0.006*** (0.000)	0.007*** (0.000)	0.006*** (0.000)
ln(1+ Listed volume USD)	0.003*** (0.001)	0.002*** (0.001)	0.003*** (0.001)	0.002*** (0.001)
Time since appraisal (days)	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
ln(1+ Competing listings volume USD)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
ln(1+ Prior-day token sales USD)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)
Annual net rent per token (USD)	0.002*** (0.001)	0.004*** (0.000)	0.002*** (0.001)	0.003*** (0.000)
Rental units	-0.000** (0.000)		-0.000** (0.000)	
Development project	0.015*** (0.002)		0.014*** (0.002)	
Renovation required	0.002 (0.002)		0.002 (0.002)	
Real estate loan	0.018*** (0.005)		0.018*** (0.005)	
Building age (years)	0.000** (0.000)		0.000** (0.000)	
Downtown distance (miles)	0.000 (0.000)		0.000 (0.000)	
Housing market	-0.000 (0.003)	0.000 (0.003)	0.001 (0.003)	0.000 (0.003)
Crypto market	0.002 (0.006)	0.007 (0.006)	0.002 (0.006)	0.007 (0.006)
Crypto market shock	-0.002 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)
1-month US Treasury yield	-0.003 (0.003)	-0.010*** (0.004)	-0.004 (0.003)	-0.010*** (0.004)
Year-month FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	No
Property FE	No	Yes	No	Yes
Observations	9,612	9,612	9,612	9,612
R^2	0.869	0.892	0.872	0.894
Within R^2	0.825	0.653	0.829	0.660

Notes: OLS on sold listings with at least one downward price revision, with dependent variable $\log(P/P^R)$, where P is the selling price and P^R is the redemption price at sale. Columns (1)–(2) use *Single-cut strategy* (=1 if exactly one cut; 0 if two or more). Columns (3)–(4) replace the binary indicator with bins of the number of cuts; the omitted category is 1 cut. Controls include total price concession (%) and premium at first cut (%), both measured as $100 \times$ the underlying fractional change (so 1 unit equals 1 percentage point). Standard errors clustered by underlying property in parentheses. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

3.7 Conclusion

Digital technologies are reshaping real-estate markets and their intermediation. Among these innovations, tokenization creates an unusually transparent setting in which listing and repricing behavior are publicly traceable. This paper applies the standard ToM framework to tokenized real estate for the first time and provides initial empirical evidence on listing behavior, seller adjustment, and marketing time for individual real-estate tokens. Using a blockchain-based peer-to-peer marketplace, we reconstruct complete listing histories from public blockchain records, including postings, price revisions, sales, and removals, together with a continuously observable valuation benchmark in the form of the redemption price. This setting allows classic questions on pricing, liquidity, and seller adjustment to be studied under a level of transparency that is typically unavailable in conventional real-estate markets.

Our results show that ToM in this setting is shaped by relative pricing and market thickness. Listings with higher premia relative to the redemption benchmark remain longer on the market, while stronger rental cash flows are associated with faster sales. A thicker stock of competing listings lengthens ToM, whereas stronger recent trading activity shortens it. Tokenization does not eliminate slow matching in real estate (Wheaton, 1990); rather, it makes ToM more tightly linked to transparent relative pricing and observable platform liquidity. The evidence further shows that traceable revision histories matter for seller adjustment. Conditional on the remaining premium, larger price reductions shorten remaining ToM, while longer revision histories are associated with more persistent illiquidity. Holding total price concessions constant, listings sold after a single decisive price cut achieve higher final premia than listings of the same underlying property repriced gradually. Buyers thus respond not only to the size of concessions, but also to how those concessions are delivered over time. These findings are especially informative because such revision-path effects are difficult to isolate in conventional real-estate settings, where posted prices are shaped by bargaining and comparable public revision paths are rarely observable. At the same time, the marketplace studied here is not a direct analogue of conventional property sales, so the mechanisms documented here should generalize most directly to other transparent digital marketplaces in which listing histories are publicly visible and execution is closely tied to posted prices. Our evidence is drawn from a single platform trading small-ticket, fractional claims on individual investment properties rather than whole properties or owner-occupied housing. Whether similar revision-path effects arise in conventional, less transparent, or more negotiation-intensive settings therefore remains open. Even so, the results suggest a broader implication for selling strategy: when price adjustments are publicly observable, repeated gradual concessions may weaken seller position, whereas a single decisive adjustment may preserve stronger execution outcomes.


More broadly, our results show that tokenization moves real-estate exchange into a transparent microstructure in which relative pricing, market thickness, and revision histories become observable and consequential for outcomes. It changes not only how individual real-estate claims trade, but also what ToM reflects. Because token markets generate high-frequency price and behavioral signals, they may also complement low-frequency transaction data and help construct token-based indicators of local market conditions, similar in spirit to Carrillo (2013), an issue for future research. As tokenization expands to other real-estate assets such as home equity and debt, future research can also examine how marketplace design and disclosure rules shape liquidity in secondary markets for those related housing-finance settings.

Appendix

A.1 Marketplace interface

Figure A.1. YAM peer-to-peer marketplace interface

- (a) Example of a real estate token listing
- (b) Listing creation interface



Note: The figure shows screenshots of the YAM peer-to-peer marketplace for real estate tokens. Panel (a) displays an example of an active secondary-market listing, including the property token, quote token, seller address, available quantity, listing price, and price/yield information. Panel (b) shows the interface for creating a new listing for an example token, where users specify the token offered for sale, the quote token, the unit sale price, and the quantity to be listed. Screenshots were taken from the YAM web interface on 4 March 2026.

A.2 Additional robustness tests

Table A.1: Determinants of Time on Market (ToM) – Removals

	Dependent variable:			
	ToM until removal (removed only) <i>OLS</i>		ToM until sample end (sales censored) <i>AFT</i>	
	(1)	(2)	<i>lognormal</i> (3)	<i>loglogistic</i> (4)
Constant	1.13 (0.95)	0.72 (0.79)	20.84*** (7.11)	14.57*** (4.55)
Listing premium (%) relative to redemption price	0.02*** (0.00)	0.02*** (0.00)	1.00*** (0.00)	1.00*** (0.00)
ln(1+ Listed volume USD)	0.10*** (0.03)	0.06** (0.03)	1.19*** (0.01)	1.17*** (0.01)
Time since appraisal (days)	-0.00 (0.00)	-0.00 (0.00)	1.00*** (0.00)	1.00*** (0.00)
ln(1+ Competing listings volume USD)	0.00 (0.01)	0.00 (0.01)	1.01* (0.00)	1.02*** (0.00)
ln(1+ Prior-day token sales USD)	-0.08*** (0.01)	-0.08*** (0.01)	0.93*** (0.00)	0.93*** (0.00)
Annual net rent per token (USD)	-0.05* (0.03)	-0.01 (0.03)	0.96*** (0.01)	0.96*** (0.01)
Rental units	0.00 (0.01)		1.00 (0.00)	1.00 (0.00)
Development project	0.22 (0.23)		1.05 (0.26)	0.93 (0.20)
Renovation required	0.12 (0.13)		1.08 (0.05)	1.02 (0.04)
Real estate loan	-0.84*** (0.29)		0.94 (0.18)	1.07 (0.18)
Building age (years)	-0.00 (0.00)		1.00*** (0.00)	1.00*** (0.00)
Downtown distance (miles)	0.01 (0.02)		1.00 (0.01)	1.00 (0.00)
Housing market	0.41** (0.16)	0.13 (0.12)	1.52*** (0.07)	1.32*** (0.06)
Crypto market	0.60 (0.40)	0.70** (0.30)	0.99 (0.22)	0.97 (0.19)
Crypto market shock	-0.41*** (0.08)	-0.35*** (0.07)	0.63*** (0.03)	0.62*** (0.02)
1-month US Treasury yield	0.01 (0.19)	0.02 (0.18)	0.78*** (0.06)	0.83*** (0.06)
Quarter-year FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	Yes
Property FE	No	Yes	No	No
Observations	35,007	35,007	111,180	111,180
R^2	0.13	0.26		
Adjusted R^2	0.13	0.24		
Log-likelihood			-142022.9	-139355.4
χ^2 (df = 79)			6363.2	6281.2

Notes: The outcome variable is ToM (T), measured as the time from listing creation until removal. All models are estimated on observations with $T > 0$. Columns (1)–(2) report OLS estimates for removed listings, with dependent variable $\log(T)$ and cluster-robust standard errors at the underlying-property level. Column (1) includes quarter-year and city fixed effects. Column (2) includes quarter-year and underlying-property fixed effects. Columns (3)–(4) report AFT models of time to removal. Listings that are sold before removal and listings still active at the end of the sample are treated as right-censored. In the AFT sample, the $T > 0$ restriction excludes 548 observations, all of which correspond to sales at $T = 0$. AFT columns report time ratios, $\exp(\beta)$, with delta-method standard errors in parentheses, where $\text{se}(\exp(\beta)) \approx \exp(\beta) \text{se}(\beta)$. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A.2: Determinants of Time on Market (ToM) – 90% sales-threshold robustness

	Dependent variable:			
	ToM until 90% sold (listings 90% sold only) <i>OLS</i>		ToM until sample end (removals censored) <i>AFT</i>	
	(1)	(2)	<i>lognormal</i> (3)	<i>loglogistic</i> (4)
Constant	-0.63 (0.65)	0.60 (0.63)	0.37*** (0.13)	0.41** (0.14)
Listing premium (%) relative to redemption price	0.14*** (0.01)	0.22*** (0.01)	1.23*** (0.00)	1.27*** (0.00)
ln(1+ Listed volume USD)	0.55*** (0.02)	0.53*** (0.02)	1.93*** (0.02)	1.92*** (0.02)
Time since appraisal (days)	-0.00*** (0.00)	-0.00*** (0.00)	1.00*** (0.00)	1.00*** (0.00)
ln(1+ Competing listings volume USD)	0.09*** (0.01)	0.09*** (0.01)	1.14*** (0.01)	1.15*** (0.01)
ln(1+ Prior-day token sales USD)	-0.55*** (0.01)	-0.49*** (0.01)	0.52*** (0.00)	0.54*** (0.00)
Annual net rent per token (USD)	-0.20*** (0.06)	-0.27*** (0.02)	0.79*** (0.01)	0.78*** (0.01)
Rental units	0.02*** (0.01)		1.03*** (0.00)	1.03*** (0.00)
Development project	-0.90 (0.85)		1.64 (0.58)	2.12** (0.74)
Renovation required	-0.04 (0.17)		0.98 (0.06)	0.93 (0.06)
Real estate loan	-1.25*** (0.47)		0.21*** (0.05)	0.21*** (0.05)
Building age (years)	-0.00 (0.00)		1.00 (0.00)	1.00*** (0.00)
Downtown distance (miles)	-0.00 (0.02)		1.02*** (0.01)	1.01* (0.01)
Housing market	-0.15** (0.07)	-0.17** (0.07)	0.86*** (0.05)	0.82*** (0.05)
Crypto market	1.22*** (0.27)	1.31*** (0.26)	2.20*** (0.63)	2.23*** (0.62)
Crypto market shock	0.23*** (0.05)	0.30*** (0.05)	1.31*** (0.07)	1.33*** (0.07)
1-month US Treasury yield	0.37*** (0.08)	0.43*** (0.08)	1.32*** (0.10)	1.25*** (0.10)
Quarter-year FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	Yes
Property FE	No	Yes	No	No
Observations	62,698	62,698	111,250	111,250
R^2	0.35	0.42		
Adjusted R^2	0.34	0.41		
Log-likelihood			-97428.9	-97071.9
χ^2 (df = 79)			40653.5	42430.7

Notes: The outcome variable is ToM (T), measured as the time from listing creation until 90% of the initially listed quantity has been sold. All models are estimated on observations with $T > 0$. Columns (1)–(2) report OLS estimates for listings that reach this 90% sales threshold, with dependent variable $\log(T)$ and cluster-robust standard errors at the underlying-property level. Column (1) includes quarter-year and city fixed effects. Column (2) includes quarter-year and underlying-property fixed effects. Columns (3)–(4) report AFT models of time to the 90% sales threshold. Listings that are removed before reaching this threshold and listings still below it at the end of the sample are treated as right-censored. In the AFT sample, the $T > 0$ restriction excludes 478 observations, all of which correspond to listings reaching the 90% sales threshold at $T = 0$. AFT columns report time ratios, $\exp(\beta)$, with delta-method standard errors in parentheses, where $\text{se}(\exp(\beta)) \approx \exp(\beta) \text{se}(\beta)$. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A.3: Revision strategy and selling price (USD)

	Dependent variable: Selling price (USD)			
	(1)	(2)	(3)	(4)
Single-revision strategy (=1)	0.010*** (0.003)	0.015*** (0.001)		
Price revisions: 2 (vs 1)			-0.009*** (0.003)	-0.011*** (0.001)
Price revisions: 3-4 (vs 1)			-0.007 (0.005)	-0.016*** (0.001)
Price revisions: 5-6 (vs 1)			-0.014** (0.007)	-0.022*** (0.003)
Price revisions: 7-11 (vs 1)			-0.028*** (0.009)	-0.034*** (0.005)
Price revisions: 12+ (vs 1)			-0.045*** (0.012)	-0.048*** (0.008)
Total price concession (%)	-0.001** (0.000)	-0.001*** (0.000)	-0.001* (0.000)	-0.001** (0.000)
Premium at first revision (%)	0.006*** (0.001)	0.006*** (0.000)	0.006*** (0.001)	0.006*** (0.000)
ln(1+ Listed volume USD)	0.005*** (0.001)	0.002*** (0.001)	0.005*** (0.001)	0.002*** (0.001)
Time since appraisal (days)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
ln(1+ Competing listings volume USD)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001** (0.000)	-0.001*** (0.000)
ln(1+ Prior-day token sales USD)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)
Annual net rent per token (USD)	0.044*** (0.013)	0.004*** (0.001)	0.044*** (0.013)	0.004*** (0.001)
Rental units	-0.000 (0.000)		-0.000 (0.000)	
Development project	0.115** (0.055)		0.115** (0.055)	
Renovation required	0.007 (0.011)		0.007 (0.011)	
Real estate loan	0.054** (0.026)		0.054** (0.026)	
Building age (years)	-0.000 (0.000)		-0.000 (0.000)	
Downtown distance (miles)	0.002 (0.001)		0.002 (0.001)	
Housing market	0.009 (0.007)	-0.001 (0.003)	0.009 (0.008)	-0.001 (0.003)
Crypto market	-0.013 (0.013)	0.002 (0.006)	-0.013 (0.013)	0.002 (0.006)
Crypto market shock	-0.003 (0.003)	0.000 (0.002)	-0.003 (0.003)	0.000 (0.002)
1-month US Treasury yield	-0.012 (0.010)	-0.006 (0.004)	-0.011 (0.010)	-0.007 (0.004)
Year-month FE	Yes	Yes	Yes	Yes
City FE	Yes	No	Yes	No
Property FE	No	Yes	No	Yes
Observations	9,612	9,612	9,612	9,612
R^2	0.722	0.950	0.723	0.951
Within R^2	0.670	0.626	0.671	0.632

Notes: OLS on sold listings with at least one downward price revision, with dependent variable $\log(P)$, where P is the selling price. Columns (1)–(2) use *Single-revision strategy* (=1 if exactly one downward price revision; 0 if two or more). Columns (3)–(4) replace the binary indicator with bins of the number of downward price revisions; the omitted category is 1 revision. Controls include total price concession (%) and premium at first revision (%), both measured as $100 \times$ the underlying fractional change (so 1 unit equals 1 percentage point). Standard errors clustered by underlying property in parentheses. Significance levels are denoted by *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

References

- Anderson K. C., Freybote J., and Manis K. T. (2024). The impact of virtual marketing strategies on the price-TOM relation. *The Journal of Real Estate Finance and Economics* 68, pp. 218–234.
- Anglin P. M., Rutherford R. C., and Springer T. M. (2003). The trade-off between the selling price of residential properties and time-on-the-market: The impact of price setting. *Journal of Real Estate Finance and Economics* 26.1, pp. 95–111.
- Baum A. (2021). Tokenization – The future of real estate investment? *The Journal of Portfolio Management* 47.10, pp. 41–61.
- BCG (2022). Relevance of on-chain asset tokenization in crypto winter. Boston Consulting Group. Accessed 2 March 2026. URL: <https://www.bcg.com/publications/2022/relevance-of-on-chain-asset-tokenization>.
- Benedetti H. R. and Rodríguez-Garnica J. P. (2023). Tokenized assets and securities. *The Emerald handbook on cryptoassets: investment opportunities and challenges*. Ed. by S. Corbet, J. W. Goodell, P. Katsiampa, and C. Miglietti. Emerald Publishing Limited, pp. 107–121.
- Bergkamp M., Sifat I., and Swinkels L. (2025). Market maturation and democratization effects of tokenized real estate. SSRN. Available at URL: <https://ssrn.com/abstract=5195172>.
- BlackRock (2025). 2025 annual chairman’s letter. Accessed 2 March 2026. URL: <https://www.blackrock.com/corporate/investor-relations/larry-fink-annual-chairmans-letter>.
- Buchak G., Matvos G., Piskorski T., and Seru A. (2025). Why is intermediating houses so difficult? Evidence from iBuyers. *Journal of Political Economy*. Forthcoming.
- Cajias M. and Freudenreich P. (2018). Exploring the determinants of liquidity with big data: Market heterogeneity in German markets. *Journal of Property Investment & Finance* 36.1, pp. 3–18.
- Carrillo P. E. (2013). To sell or not to sell: Measuring the heat of the housing market. *Real Estate Economics* 41.2, pp. 310–346.
- Carrillo P. E. and Williams B. (2019). The repeat time-on-the-market index. *Journal of Urban Economics* 112, pp. 33–49.
- Cornelli G. (2025). When bricks meet bytes: Does tokenisation fill gaps in traditional real estate markets? BIS Working Papers 1311. Bank for International Settlements.
- Deloitte (2025). Digital dividends: How tokenized real estate could revolutionize asset management. Deloitte Insights. Accessed 2 March 2026. URL: <https://www.deloitte.com/us/en/insights/industry/financial-services/financial-services-industry-predictions/2025/tokenized-real-estate.html>.
- Eerola E. and Lyytikäinen T. (2015). On the role of public price information in housing markets. *Regional Science and Urban Economics* 53, pp. 74–84.
- Ferreira E. J. and Sirmans G. S. (1989). Selling price, financing premiums, and days on the market. *Journal of Real Estate Finance and Economics* 2, pp. 209–222.
- Gedikli C., Hill R., Talavera O., and Yilmaz O. (2026). Online real estate agencies and their impact on the housing market. *Real Estate Economics* 54.2, pp. 540–579.
- Genesove D. and Han L. (2012). Search and matching in the housing market. *Journal of Urban Economics* 72, pp. 31–45.

- Glascok J. L., Lu C.-L., and So R. W. (2000). Further evidence on the integration of REIT, bond, and stock returns. *The Journal of Real Estate Finance and Economics* 20.2, pp. 177–194.
- Glower M., Haurin D. R., and Hendershott P. H. (1998). Selling time and selling price: The influence of seller motivation. *Real Estate Economics* 26.4, pp. 719–740.
- Haurin D. R. (1988). The duration of marketing time of residential housing. *AREUEA Journal* 16.4, pp. 396–410.
- Haurin D. R., Haurin J. L., Nadauld T., and Sanders A. (2010). List prices, sale prices and marketing time: An application to U.S. housing markets. *Real Estate Economics* 38.4, pp. 659–685.
- Herrin W. E., Knight J. R., and Sirmans C. F. (2004). Price cutting behavior in residential markets. *Journal of Housing Economics* 13, pp. 195–207.
- Horowitz J. L. (1992). The role of the list price in housing markets: Theory and an econometric model. *Journal of Applied Econometrics* 7.2, pp. 115–129.
- Kalbfleisch J. D. and Prentice R. L. (2002). *The statistical analysis of failure time data*. 2nd ed. Wiley-Interscience.
- Kalyani A., Bloom N., Carvalho M., Hassan T., Lerner J., and Tahoun A. (May 2025). The diffusion of new technologies. *The Quarterly Journal of Economics* 140.2, pp. 1299–1365.
- Kang H. B. and Gardner M. J. (1989). Selling price and marketing time in the residential real estate market. *The Journal of Real Estate Research* 4.1, pp. 21–35.
- Knight J. R. (2002). Listing price, time on market, and ultimate selling price: Causes and effects of listing price changes. *Real Estate Economics* 30.2, pp. 213–237.
- Konashevych O. (2020). General concept of real estate tokenization on blockchain: The right to choose. *European Property Law Journal* 9.1, pp. 21–66.
- Kreppmeier J., Laschinger R., Steininger B. I., and Dorfleitner G. (2023). Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *Journal of Banking and Finance* 154, p. 106940.
- Laschinger R., Leonhard H., Dorfleitner G., and Schäfers W. (2026). Liquidity mechanisms in decentralized finance: Design and fragmentation in illiquid real-world asset markets. Working Paper.
- Lazear E. P. (1986). Retail pricing and clearance sales. *The American Economic Review* 76.1, pp. 14–32. URL: <https://www.jstor.org/stable/1804124>.
- Ling D. C. and Naranjo A. (1999). The integration of commercial real estate markets and stock markets. *Real Estate Economics* 27.3, pp. 483–515.
- Liu A. and Chen C. (2025). From real estate financialization to decentralization: A comparative review of REITs and blockchain-based tokenization. *Geoforum* 159, p. 104193.
- Lukkarinen A. and Schwienbacher A. (2023). Secondary market listings in equity crowdfunding: The missing link? *Research Policy* 52.1, p. 104648.
- Merlo A. and Ortalo-Magné F. (2004). Bargaining over residential real estate: Evidence from England. *Journal of Urban Economics* 56, pp. 192–216.
- Merlo A., Ortalo-Magné F., and Rust J. (2015). The home selling problem: Theory and evidence. *International Economic Review* 56.2, pp. 457–484.
- Miller N. G. (1978). Time on the market and selling price. *AREUEA Journal* 6.2, pp. 164–174.

- Miller N. G. and Sklarz M. A. (1987). Pricing strategies and residential property selling prices. *The Journal of Real Estate Research* 2.1, pp. 31–40.
- Read C. (1988). Price strategies for idiosyncratic goods: The case of housing. *AREUEA Journal* 16.4, pp. 370–395.
- Sass T. R. (1988). A note on optimal price cutting behavior under demand uncertainty. *The Review of Economics and Statistics* 70.2, pp. 336–339.
- Schär F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review* 103.2, pp. 153–174.
- Sing T. F. and Zou Y. (2024). Improving search efficiency in housing markets using online listings. *Real Estate Economics* 52.4, pp. 1020–1044.
- Smith J., Vora M., Benedetti H., Yoshida K., and Vogel Z. (2019). Tokenized securities & commercial real estate. MIT Digital Currency Initiative. SSRN. Available at URL: <https://ssrn.com/abstract=3438286>.
- Steininger B. I. (2023). Return–risk analysis of real estate tokens: An asset class of its own. *The Journal of Portfolio Management* 49.10, pp. 83–102.
- Swinkels L. (2023). Empirical evidence on the ownership and liquidity of real estate tokens. *Financial Innovation* 9, pp. 1–29.
- Taylor C. R. (1999). Time-on-the-market as a sign of quality. *The Review of Economic Studies* 66.3, pp. 555–578.
- Trippi R. R. (1977). Estimating the relationship between price and time to sale for investment property. *Management Science* 23.8, pp. 838–842.
- Verstappen J. (2025). Tokenisation of Absolute Rights and Claims: On the Use of Tokens to Transfer Rights in Property Law. *European Journal of Risk Regulation*. FirstView, pp. 1–20.
- Wei L.-J. (1992). The accelerated failure time model: A useful alternative to the Cox regression model in survival analysis. *Statistics in Medicine* 11.14–15, pp. 1871–1879.
- Wheaton W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy* 98.6, pp. 1270–1292.
- Wit E. R. de and Klaauw B. van der (2013). Asymmetric information and list-price reductions in the housing market. *Regional Science and Urban Economics* 43, pp. 507–520.
- Yavas A. and Yang S. (1995). The strategic role of listing price in marketing real estate: Theory and evidence. *Real Estate Economics* 23.3, pp. 347–368.

Conflicts of Interest

One author bought and sold few digital tokens issued by the company RealT so that they could retrace the process and trading of tokenization. The current value of these holdings is less than \$100.

Chapter 4

Virtual Land in the Metaverse? Exploring the Dynamic Correlation with Physical Real Estate

This chapter is based on a co-authored paper with Maximilian Nagl (Chair of Statistics and Risk Management, University of Regensburg) and Wolfgang Schäfers (IREBS, International Real Estate Business School, University of Regensburg). It has been published as Leonhard, H., Nagl, M., and Schäfers, W. (2025), “Virtual land in the metaverse? Exploring the dynamic correlation with physical real estate,” *Journal of European Real Estate Research*, 18(1), 121–151, as part of the Special Issue on Digital Technology and Innovation in Real Estate.

Abstract

Purpose: As blockchain-based virtual worlds gain prominence within the emerging metaverse and Web3, numerous global companies and investors are buying purely virtual land to explore new business potentials and capitalize on digital assets. Given the similarities to physical real estate, this study examines the dynamics of the secondary market for virtual land and relates its returns to those of physical real estate.

Design/methodology/approach: Using transaction-level data from a prominent virtual land platform, the authors construct a virtual land market index based on repeat sales index methodology from traditional real estate studies. Wavelet coherence analysis is employed to examine the dynamic correlation between virtual land and various physical real estate market returns. The determinants of this correlation are estimated using stepwise regression analysis. A portfolio analysis explores the implications of adding virtual land to traditional asset portfolios.

Findings: The correlation between virtual and physical real estate market returns is generally low, reaching its lowest during the Covid-19 lockdowns from 2020 to 2022. It spikes during acute economic turmoil such as the initial Covid-19 outbreak or interest rate change announcements. The correlation is primarily driven by consumer and economic climate, the price of the virtual economy token and investor attention. Portfolio analysis indicates that virtual land can enhance risk-adjusted returns within a traditional portfolio, particularly when added to a commercial real estate portfolio.

Research limitations/implications: This study examines a single virtual land market, despite it being the oldest and one of the largest. Given the rapidly evolving nature of virtual worlds, it is crucial to further test the results and include new virtual land platforms as they emerge.

Practical implications: The findings provide actionable insights on portfolio implications for investors seeking alternative real-estate-like assets in the digital space. Additionally, this study offers strategic guidance for entering the metaverse, including a comprehensive overview of established virtual presences.

Originality/value: With the advancing digitization of real estate markets, this study is the first to explore the correlation between market returns of virtual land in the metaverse and traditional physical real estate. The findings provide valuable empirical insights for investors, policymakers, entrepreneurs and companies interested in the intersection of digital and traditional property markets.

Keywords: Digital Assets, Blockchain, Metaverse, Virtual Land, Real Estate, Wavelets

4.1 Introduction

Blockchain technology has introduced significant advancements and transformative potential across various industries. In particular, its application in real estate markets has witnessed notable developments, such as the tokenization of physical properties (Kreppmeier et al., 2023; Swinkels, 2023; Simons and Simons, 2022) or the emergence of blockchain-based real estate firms (Nagl et al., 2024). By leveraging the unique features of blockchain, investors find new opportunities to include new assets into their portfolio.¹ A rather novel innovation is the creation of solely digital real estate, known as virtual land, which refers to digitally owned, blockchain-based parcels of land in virtual worlds or metaverse platforms. Metaverse platforms have witnessed a surge in commercial activities, attracting substantial investments, time, attention and corporate engagement.² For example, Deutsche Bank established its metaverse presence in August 2023. It can be explored at <https://play.decentraland.org/?position=-64%2C1> (accessed on April 13th, 2024). A virtual land parcel is represented by a Non-Fungible Token (NFT), a unique, indivisible, irreplaceable, verifiable, and tradable token that represents ownership and proof of authenticity for a specific digital or physical object, such as (virtual) real estate, art, or content and is stored on a blockchain.

The metaverse envisions a contiguous spatial digital universe, comprising a network of interoperable and navigable platforms and spaces. It can be analogized to a quasi-geographical, 3D cityscape World Wide Web (Schumacher, 2022). Thus, the idea of owning virtual land can be compared to owning an internet domain name (Lindenthal, 2018). Similar to physical urban economics, the more popular a metaverse platform (or city) gets, the higher the attention and the price of its terrain (Schär and Goldberg, 2023). While the concept of virtual economies originated in the internet gaming industry in the early 2000s (Knowles and Castronova, 2016), the integration of decentralized blockchain infrastructure, often referred to as Web3, has revolutionized the field by diminishing reliance on central operators, who could, with a flip of a switch, alter or shut down their virtual platforms.

In this paper, we compile and analyze a comprehensive transaction and parcel data set from the popular blockchain-based metaverse platform *Decentraland* (Ordano et al., 2017) and investigate the correlation between its virtual land market returns and physical real estate market returns. In the following, we introduce virtual land and compare it to its physical counterpart.

Valuation Both serve as valuable assets, with physical real estate representing tangible properties in the real world and virtual land representing ownership rights in the digital space. Similar to physical real estate, valuation of virtual land in the metaverse is driven by location and expected visitor density (Goldberg et al., 2024; Yench, 2023), demand-supply dynamics (Momtaz, 2022), and market conditions (Vidal-Tomás, 2022; Maouchi et al., 2022). Virtual land valuation integrates elements of internet domain markets, where easy-to-remember and special domain names are positively correlated to their prices (Lindenthal, 2014). In the case of virtual land platforms, price premiums can be observed for coordinates that feature repeating digits (Goldberg et al., 2024). Similar numerological phenomena can also be observed in physical real estate markets (see, e.g., Humphreys et al., 2019).

¹In the place of many, see “HSBC launches metaverse portfolio for wealthy Asian clients” <https://www.reuters.com/technology/hsbc-launches-metaverse-portfolio-wealthy-asian-clients-2022-04-06/>.

²See for example <https://www.bbc.com/news/technology-63488059>.

Transactions Both, virtual and physical land can be bought, sold, rented and transferred on an existing market, although virtual land transactions are processed on blockchain-based platforms utilizing smart contracts to ensure maximum transparency, security and immutability. The current total virtual land market of \$2.5 billion (OneLand, 2024) is still only a fraction of the compared \$338 trillion physical real estate market (Savills, 2023). However, the metaverse asset market is expected to experience explosive growth in valuation in the coming years, emphasizing the need to demystify its entanglement with broader asset markets.³

Real Estate Types In addition to their similarities, physical real estate offers physical spaces for various activities, including residential, commercial, industrial or agricultural purposes. On the other hand, virtual land supports diverse applications mainly of commercial purpose such as virtual storefronts, events (Goldberg et al., 2023), communities, gaming environments (Proelss et al., 2023), marketplaces, and other decentralized applications (DApps). We find virtual locational choice to resemble (commercial) real-world preferences concentrating along matching districts and clusters of similar companies, roads, and plazas. Both asset classes offer opportunities for development, and the stimulation of creativity as well as entrepreneurship and economic activity.

Land Scarcity Importantly, both physical and virtual land incorporate scarcity characteristics, making them valuable assets that cannot be duplicated arbitrarily. Physical real estate supply is constrained by the earth's surface area and the principles of gravity. Blockchain-based virtual land supply is determined by a decentralized autonomous organization (DAO), which is the community-led governance entity with no central leadership. Voting-right is based on the participants stake in the platform in form of virtual land or platform economy tokens (Goldberg and Schär, 2023).

Physical vs Digital Despite these similarities, there are notable differences between virtual and physical real estate. The most significant distinction is the physical nature of traditional real estate, which can be experienced, occupied, and utilized in the physical world. In contrast, virtual land exists solely in the digital space and lacks a physical presence and tangibility. This unchains new unexperienced spontaneity, flexibility, and architectural creativity of virtual land development since the construction is based on bytes and bits rather than bricks and mortar (Schumacher, 2022).

Regulation Physical real estate operates within well-established legal and regulatory frameworks, while virtual land operates within a relatively nascent and evolving regulatory landscape (Goanta, 2020). The buying and selling process for virtual land is generally cheaper, faster and (yet) without the tremendous tax implications physical real estate transactions face.

Data Availability Researchers also acknowledge another important distinction of virtual land compared to physical real estate markets: the availability of public blockchain data, which is freely accessible, while physical real estate data is often opaque, leading to delays and challenges. This allows not only for transparent transaction data but also emphasizes the benefits of using blockchain-based virtual worlds to conduct real-world policy and urban planning experiments (see, e.g., Saengchote et al., 2023; Hudson-Smith, 2022).

Despite the interesting similarities between virtual land in the metaverse and its physical

³See <https://www.statista.com/outlook/amo/metaverse/metaverse-virtual-assets/worldwide>.

equivalent, the motivation for this study arises from its novelty and the limited research conducted so far. One of the first empirical studies in this area, conducted by Wang et al. (2010), suggests a potential influence of physical real estate prices on virtual land prices in Second Life, a first-generation non-blockchain-based virtual world.⁴ However, the findings are constrained by a limited methodology, a dataset covering only one year, and, most importantly, the study was conducted before the emergence of blockchain-based ownership and NFTs. In the context of asset pricing studies and interrelatedness investigations of virtual land with other asset classes, virtual land has either been aggregated with other NFTs (mostly digital art) (see, e.g., Aharon and Demir, 2022; Ante, 2022; Ko and Lee, 2023; Umar et al., 2022; Yousaf and Yarovaya, 2022) or real estate as a traditional asset class was not included into the comparisons (see, e.g., Ko et al., 2022; Urom et al., 2022; Dowling, 2022; Kumar and Padakandla, 2023). In summary, all studies found low interaction between NFTs and traditional asset classes.

Our paper aims to explicitly investigate the relationship between virtual land and physical real estate and address the observed research gap. Therefore, we aim at answering the following research questions:

- *First*, are virtual land market returns connected to physical land market returns?
- *Second*, does the connectedness vary for different investment horizons and time frames?
- *Third*, which drivers can explain the connectedness of virtual and physical land?

We analyze the returns of the largest blockchain-based virtual land market, Decentraland, and a broad compilation of physical commercial real estate in form of a REIT index. Our data reaches from 2019 until 2023 which is, to the best of our knowledge, the longest time frame observed for the performance of virtual land. We adopt repeat sales index methodology from traditional real estate studies proposed by Case and Shiller (1987) to construct an aggregated virtual land repeat sales index. Due to the predominately commercial setting of the observed virtual land market, we isolate the retail and office exposure from the well-known FTSE Nareit All Equity REITs Index. Since REITs are frequently traded like virtual land, we ensure data availability and comparability. We analyze the dynamic correlation between the two market returns by employing wavelet coherence analysis.

We discover that the correlation between virtual and physical land returns is generally low, reaching its nadir for medium-term investment horizons spanning two to eight months and particularly during the Covid-19-induced lockdown period in 2020 until 2022. Conversely, the correlation peaks during short-term extreme events, such as the acute Covid-19 outbreak in early 2020, announcements of interest rate changes, and brief periods of stock market volatility lasting up to two months, similar to patterns observed in equity REITs. Our findings suggest that the primary drivers of correlation between both market returns include the consumer and economic climate, the price of Decentraland's cryptocurrency, MANA, and investor attention to both asset types.

It is important to note that our study is just a starting point for exploring the fast-changing virtual real estate landscape. However, our results provide initial insights for market participants, policymakers, and researchers. First, the insights offer guidance for investors

⁴See <https://secondlife.com/>.

considering virtual land market investments. In this case, our results suggest strategic diversification and portfolio hedging effects, especially for retail and commercial-focused real estate investors aiming for alternative real-estate-like assets in the digital realm. Based on our results, virtual land has the potential to enhance risk-adjusted returns within a traditional asset portfolio. Companies investing in virtual land to extend their business operations can gain foundational knowledge about virtual land markets from a real estate perspective and benefit from our unique and hand-sourced list of companies already engaging in the metaverse. Furthermore, our study can be of interest to PropTech and Web3 entrepreneurs aiming to launch new virtual land platforms and capitalize on the digitization of real estate markets. Ultimately, our findings can help policymakers take well-informed regulatory actions, better categorize virtual land as an emerging asset class, and do so without hindering innovation.

The remainder of this paper is structured as follows. In Section 4.2, we describe the commercial setting of virtual land markets and reveal insights about its corporate presences. In Section 4.3 we describe our data and in Section 4.4 we introduce the empirical methods. The results are presented in Section 4.5. In Section 4.6 we conclude.

4.2 Virtual land market

Decentraland is the first large-scale blockchain-based virtual world, which enables browser-based access since February 20, 2020. Users can navigate as avatars and purchase, develop, and trade virtual land and assets as well as socialize, explore, shop, and play games (Ordano et al., 2017). In contrast to its largest competitor *The Sandbox* where movement is restricted to jumping from one parcel to another, Decentraland offers a seamless open-world environment that allows users to roam freely between parcels. Since the creation of Decentraland, its surface is divided into equal-sized 92,598 parcels, which can be classified into four distinct categories: districts (35,883), private land parcels (43,689), public plazas (3,588), and roads (9,438). Districts vary in size, location, and themes and encompass activities such as shopping, gambling, education, art, music, and business. While the road system was originally intended for transportation, users typically rely on teleportation to move between different areas. However, its utility as a spatial browsing tool explains the increased demand and price premiums for adjacent parcels (Goldberg et al., 2024). There are nine plazas in Decentraland, symmetrically arranged around the central plaza, resembling a polycentric city model. The central plaza serves as an entry point to Decentraland, accessible by logging into the virtual world. The virtual world is governed by the community-owned Decentraland DAO, which serves as the hub for decision-making and establishing the rules of the virtual world. Users have the opportunity to participate in proposals and exercise their voting rights based on their virtual land plots and platform tokens, as further investigated by Goldberg and Schär (2023).

All private and district parcels were initially sold at public auctions in 2017 and 2018. Goldberg et al. (2024) and Yenchu (2023) analyzed these auction prices and found significant locational effects, despite the negligible transportation costs in the virtual world. Once sold on the primary market, virtual land can only be bought and traded on secondary marketplaces. It is possible to connect adjacent land parcels to create larger scenes and maximize content discoverability, a practice that is common and lucrative in virtual real estate markets, as

concluded by Yi (2019).⁵

Ante et al. (2023) find that retail investor motivation for virtual land ownership is driven by intrinsic factors such as perceived aesthetics and digital self-expression, as well as extrinsic factors like potential financial gains from speculation, investment, and the appeal of investing in technological innovations and innovative business models. The latter is likely to influence corporates as well when considering the purchase of virtual land.⁶ Furthermore, companies buy virtual land to deepen customer relationships by engaging with a younger, digital-native audience that increasingly spends more time in virtual spaces and largely considers online experiences to be as meaningful as in-person interactions (Westcott et al., 2023). Corporate metaverse presences can enable more personalized and gamified brand and community experiences and create a new way of advertising (Giang Barrera and Shah, 2023; Hollensen et al., 2022). For example, fashion clothing is promoted in the metaverse during retail events and can be purchased for both the physical world and as digital wearables for personal avatars in the virtual world (Goldberg et al., 2023). Jain et al. (2022) showcases a broad selection of retail sale activities in the metaverse and describes it as an e-commerce revolution.

We find 76 major global companies and institutions across different sectors and countries with a virtual presence in Decentraland (Figure 4.1).⁷ Similar to real-world commercial real estate markets, the mapped corporate locations reveal preferences for proximity to matching districts, similar companies, roads, and plazas, following well-known real-estate principles of agglomeration economies, co-location synergies, and network effects (Schumacher, 2022). In particular, we find two clusters around *Crypto Valley* and *Fashion Street* with luxury retailers like Dolce & Gabbana and Balenciaga and mass-market retailers like Adidas and Tommy Hilfiger (Figure 4.2). Besides fashion brands, we discovered 16 banking & financial corporations including Banco de Chile, Fidelity, Oversea-Chinese Banking Corporation, and Arab National Bank.⁸ Banking in the metaverse is primarily about reaching younger and crypto-native customers, creating an innovative image of the company and simply experimenting with the new technologies involved (Ooi et al., 2023). We find a few real estate and construction firms active in Decentraland, notably Swiss Prime Site, which built a digital replica of their Prime Tower in Zurich that can be explored in the virtual world.⁹

⁵Scenes on virtual land plots include 3D objects and audio content, created either independently or by metaverse architects like <https://landvault.io/>. To optimize performance and manage land scarcity, scene design follows building height restrictions based on parcel count, starting at 20m for a single parcel. Scenes exceeding these limits are not loaded.


⁶See, e.g., Forbes (2024): “Our venture into [...] metaverse is a testament to Forbes’ commitment to embracing the forefront of technological innovation,” said Taha Ahmed, Forbes’ Chief Growth Officer (<https://www.forbes.com/sites/digital-assets/2024/02/19/forbes-launches-permanent-presence-in-the-sandbox-metaverse/>).

⁷For a complete list of companies, refer to Table A.1 in the Appendix. For examples of corporate metaverse presences, see Figure A.1 also in the Appendix.

⁸For a complete breakdown of companies by business sectors and continents, refer to Table A.2 and A.3 in the Appendix.

⁹See <https://sps.swiss/en/group/real-estate/portfolio/office/prime-tower-1>.

Figure 4.1: Decentraland map and companies (n=76)



Source(s): Authors' own work based on data from Decentraland.


4.3 Data

4.3.1 Virtual land transactions

Now that the observed virtual land market and its commercial context has been outlined, we will present the transaction data used in our study. For our analysis, we exclusively focus on Decentraland, the only blockchain-based large-scale virtual land platform that provides an open-world environment. In Decentraland, avatars can move seamlessly across parcels without boundaries, enabling spontaneous exploration and much closely resembling the dynamics of a real-life city and property market. This stands in contrast to another very popular blockchain-based virtual land platform, *The Sandbox*, where participants can only teleport between parcels. This limitation reduces the possibility of spontaneous encounters with other virtual presences, potentially leading to different market dynamics.¹⁰ We retrieve all transactions of virtual land parcels in Decentraland, so called *LAND*, between January 1, 2019 and December 31, 2023, which amounts to 18,212 sales. We collect the data from *OpenSea*, the

¹⁰Furthermore, Decentraland is the oldest blockchain-based virtual land market, with all parcels sold to the public in December 2018. This early adoption has resulted in a richer dataset for the secondary market, despite the relatively short period analyzed. In contrast, *The Sandbox*, for instance, only began its first public land sale in February 2021, resulting in a much shorter dataset.

Figure 4.2: Fashion district



Source(s): Authors' own work based on data from Decentraland.

first and largest NFT marketplace by transactions and sales volume.¹¹ LAND can be listed in different cryptocurrencies, including MANA, Ether and Wrapped Ether. We convert all denominations to their corresponding U.S. dollar equivalents based on the daily rates from the leading cryptocurrency database CoinMarketCap (2024).¹² Each transaction contains specific information, including the transaction price, the buyer and seller addresses, the timestamp, the unique global identification code (token ID), which allows for the identification of resales and meta data like the Cartesian coordinates. To ensure a homogeneous dataset for the index construction, we exclude 1,873 transactions of adjacent connected parcels (Estates), 89 transactions of multiple parcels which are not necessarily adjacent (bundles), and 395 private sales, whose arrangement happens over-the-counter. Ultimately, we obtain 15,855 non-private transactions of 7,665 unique single parcels during January 1, 2019 and December 31, 2023.

For the repeat sales index construction, we aggregate our data into weekly intervals to ensure sufficient transactions per time period and to reduce the number of potential non-arm's length transactions and wash trades, which are indicated by short time periods in between the sales and are commonly encountered in NFT markets (Serneels, 2023).

4.3.2 Real estate markets

We use publicly traded real estate assets in form of real estate investment trusts (REIT) to approximate the returns of the physical real estate market. We focus on equity REITs, since virtual land is (yet) predominantly bought with equity. REITs are traded on a daily basis and thus guarantee a sufficient number of observations. Most importantly, the price discovery in

¹¹Decentraland is built on the Ethereum blockchain (Buterin, 2014; Wood, 2014) and LAND tokens follow the Ethereum Request for Comments 721 (ERC-721) token standard. Opensea provides a freely accessible application programming interface (API) to access its NFT data, which is compiled and structured from the Ethereum blockchain <https://docs.opensea.io/reference/api-overview>.

¹²72 % of our observed transactions are settled with MANA, the native economy token of Decentraland, followed by Ether, the native token of the underlying Ethereum blockchain.

the REIT markets usually precedes that in the private markets and both markets tend to co-move in the long run (see, e.g., Hoesli and Oikarinen, 2021; Hoesli and Oikarinen, 2016; Hoesli and Oikarinen, 2012; Boudry et al., 2012; Shaver et al., 2023; Pagliari et al., 2005; Clayton and MacKinnon, 2003).

We follow the approach by Nagl et al. (2024), which utilize the FTSE Nareit All Equity REITs Index to examine the return dependence between real estate crypto coins and real estate markets. For an adequate comparison in our analysis, we further modify the index according to the observed predominant commercial setting in the virtual land market and exclude all real estate sectors other than retail, office, and gaming.¹³ The sectoral weighting of our sub-index is based on the quarterly weighting of the reference index, which on average is 63% retail, 35% office, and 3% gaming. To enhance our analysis, we also compare the virtual land market returns with the returns of the entire FTSE Nareit All Equity REITs Index and the worldwide S&P Global REIT Index.

4.4 Methodology

4.4.1 Repeat sales index

To benchmark the performance of the virtual land market and conduct correlation and portfolio analyses with physical real estate, we construct a repeat sales index of parcel-level virtual land sales. We apply traditional real estate index methodology by Case and Shiller (1987), which handles the problem of individual properties transacting irregularly and infrequently by estimating the missing periodic values using a generalized-least square regression (GLS), which is based on only repeat sales transactions.¹⁴ In addition to its real estate application (see, e.g., S&P/Case-Shiller Home Price Index), the repeat sales methodology is often used for other irregularly traded and non-standardized assets like corporate bonds (Beaupain and Heck, 2016), art (Pesando, 1993; Mei and Moses, 2002), Internet domain names (Lindenthal, 2014), and NFTs (Kong and Lin, 2021; Borri et al., 2023; Nakavachara and Saengchote, 2022).

The approach assumes constant characteristics of the repeatedly sold properties, which is often criticized in the context of real estate. Haurin and Hendershott (1991, p. 260) summarize the disadvantages of the repeat sales method as follows: *“(1) it does not separate house price change from depreciation, (2) renovation between sales is ignored, (3) the sample is not representative of the stock of housing, (4) attribute prices may change over time, and (5) a large number of sales are required before a reasonable repeat-sales sample is obtained.”*

Yet, when applied to the virtual land market of Decentraland, these disadvantages can be mitigated: (1) There is no kind of physical depreciation of virtual land, as it is digital and (2) renovations and repairments are negligible. (3) Most of the traded virtual land was repeatedly sold and the sample used for index construction should be representative. (4) The digital nature diminishes the issues of attribute price changes like labor, capital, and materials to an

¹³As analyzed by Schär and Goldberg (2023), participants in virtual worlds have no need for housing, food, or other essentials, resulting in minimal demand for residential, hotel, agricultural, or industrial land in the metaverse.

¹⁴The repeat sales methodology is explained in more detail in Appendix A.3.

irrelevant extent. (5) We find a large enough number of sales, since the change of hands of virtual land is higher compared to traditional real estate. More than half of the observed parcels are repeatedly sold. In addition to these mitigations, it is appealing that all plots of Decentraland have the same size (16 m x 16 m) and the same building restrictions.

4.4.2 Wavelet analysis

Understanding the co-movement between financial assets is crucial for risk management, portfolio optimization, and forecasting. Wavelet analysis is a versatile mathematical tool for analyzing and extracting information from non-stationary time series data (Torrence and Compo, 1998). This is particularly valuable in finance, where non-stationary time series frequently occur. Furthermore, while standard approaches such as dynamic conditional correlation (DCC) or Vector Autoregression (VAR) allow us to study the covariance solely in the time domain, the wavelet analysis additionally uncovers the frequency domain. In the context of financial time series data the frequency can be interpreted as the investment horizon. The application of bivariate wavelet analysis to study the coherence of two financial time-series is well approbated (see, e.g., Rua and Nunes, 2009 for stock returns and Dowling, 2022; Umar et al., 2022; Vidal-Tomás, 2022 for NFT and cryptocurrency returns). In our study, we analyze the coherence between virtual land and physical real estate returns.

We calculate the wavelet squared coherency, which measures the similarity or coherence to which two time series move together around each moment in time and for each frequency (time–frequency domain). Its intuition can be interpreted like the squared coefficient of correlation, since it also ranges between 0 (no co-movement/correlation) and 1 (perfect co-movement/correlation). But, the wavelet approach allows for the detection of both time-varying and frequency-varying features, providing a more comprehensive understanding of the relationship between the two time series. Following Torrence and Compo (1998) and Grinsted et al. (2004), we calculate the cross-wavelet spectrum by taking the absolute value squared of the smoothed cross-wavelet spectrum $W_{xy}(\tau, s)$ and normalizing it by the product of the absolute value squared of the smoothed wavelet transforms $W_x(\tau, s)$ and $W_y(\tau, s)$ as follows¹⁵

$$R^2(\tau, s) = \frac{|S(s^{-1}W_{xy}(\tau, s))|^2}{S(s^{-1}|W_x(\tau, s)|^2) S(s^{-1}|W_y(\tau, s)|^2)}. \quad (4.1)$$

By analyzing the graph of the wavelet squared coherency (scalogram), one can identify the specific regions in the time-frequency domain where the two time series exhibit co-movement.

4.5 Results

4.5.1 Virtual land transactions

We begin the results section by visualizing and examining the underlying virtual land transaction data which sets the basis for our index construction in the next subsection.


¹⁵For a more detailed derivation of the formula, please see Appendix A.4.

This subsection aims to identify key factors that influence pricing dynamics at the parcel level before aggregating the data into a market index. The combined temporal and spatial representation of virtual land sales helps to better understand virtual land market dynamics within a real estate-like context. For a more in-depth analysis of transaction activity over time and investor participation, further figures and tables can be found in Appendix A.2.


4.5.1.1 Transaction prices per parcel

First, we present the cadastral distribution of the average transaction prices in Figure 4.3 and 4.4. The maps reveal that maximum parcel prices can be located around the central plaza, along roads and intersections in proximity to plazas as well as established agglomeration like the previously highlighted *Fashion Street* and *Crypto Valley*. For example, Figure 4.4 reveals that the few parcels with average prices over \$100,000 are adjacent or in closest proximity to the central plaza. The center-of-gravity effect comes with no surprise since the users begin their virtual journey always at the central plaza, when logging into the platform and thus pedestrian traffic and content discoverability is greatest. A similar effect is given along the road system, which besides teleportation is the intuitive way to browse the virtual world.

Figure 4.3: Average transaction prices ($n = 7,665$)



Source(s): Authors' own work based on data from Opensea and Decentraland.

Figure 4.4: Average transactions prices $\geq \$10,000$ ($n = 2,182$) – Enlarged around the center

Source(s): Authors' own work based on data from Opensea and Decentraland.

Our results for the secondary market of virtual land confirm the findings of previous studies investigating virtual land prices on the primary market (Goldberg et al., 2024; Yench, 2023). The findings for virtual land transactions emphasize the comparable locational price dynamics experienced in real-world commercial and retail real estate markets, where high visibility and foot traffic as well as central positions correlates with higher price premiums (see, e.g., Heikkilä et al., 1989).


After presenting the spatial distribution, we present the temporal distribution of all observed transaction prices in Figure 4.5 and show the corresponding summary statistics in Table 4.1. The following key takeaways can be extracted. Firstly, during periods of high transaction activity represented by the dark concentration of dots especially in Q2 2021 and between Q4 2021 and Q1 2022, the range of transaction prices expands significantly. Secondly, there are high temporal differences in average transaction prices ranging between \$771 in 2019 to \$10,542 in 2021. The phase of increased prices and trading activity in 2021 and 2022 coincides with the increased media attention around virtual land and the metaverse.¹⁶ Due to the consequences of Covid-19 at this time, including social containment, travel restrictions and lockdowns, the demand for virtual realms to socially interact, work, shop and play surged. Furthermore, the introduction of NFTs made decentralized ownership of these digital products and content in virtual worlds possible. However, in 2023, we can observe a significant decrease of virtual land transactions and prices as the market adjusted to unmet expectations surrounding the metaverse that have yet to be realized along with the current lack of mass adoption. Despite this reduction in market and transaction activity, 2023 still saw some notable virtual land transactions and marketing activities. For instance, Deutsche Bank purchased virtual land for an expanded presence in Decentraland on August 22, and conducted its first Metaverse event on December 14, where the Chief Investment Officer presented his 2024 outlook as a live avatar.¹⁷ Other notable acquisitions of virtual land in

¹⁶In the place of many, see <https://www.cnn.com/2022/01/12/investors-are-paying-millions-for-virtual-land-in-the-metaverse.html>.

¹⁷For more details, see <https://www.db.com/what-next/digital-disruption/Metaverse/banking/> and

Decentraland in 2023 include Banco de Chile on June 16, Nestlé on April 27, and Fidelity on March 17.¹⁸ Additionally, many fashion brands maintain their virtual presence, that was activated during the Metaverse Fashion Week in Decentraland in March 2023 and that took place in the Fashion District (see Figure 4.2).¹⁹ To further underline the ongoing market activity, it is worth noting that the average transaction price in 2023 (\$1,433) remained double that of 2019 (\$771). To some extent, comparable cyclical fluctuations can also be observed in physical real estate markets, particularly in the commercial and office property sectors (see, e.g., DiPasquale and Wheaton, 1992). Even though the underlying factors for these cyclical patterns may differ substantially, one being driven by intangible aspects and the other by tangible ones, both virtual and physical real estate provide functional space that is subject to demand and supply shifts, popularity, location, and speculative interests. These cycles can present both challenges and opportunities for investors, depending on their timing and strategy. Investors with a longer investment horizon may view the current lower market phase and reduced transaction prices as an opportunity to acquire assets at a discount, with the expectation of future appreciation. Similarly, companies may be attracted by lower virtual land prices to start experimenting within these metaverse platforms.

Figure 4.5: Transaction prices (n = 15,855)



Source(s): Authors' own work based on data from Opensea.

<https://www.instagram.com/deutschebank/reel/C29VQH0KxPa/>.

¹⁸For Banco de Chile, see <https://portales.bancochile.cl/personas/dimensionb>. For Nestlé, see <https://www.nestle-family.com/en/nestle-cerealmclub>. For Fidelity, see <https://www.fidelityinternational.com/about-us/fidelity-dcl-experience-activation/>.

¹⁹For more details, see <https://mvfw.org/>.

Table 4.1: Transaction prices – Summary statistics

Year	N	Mean	Std. Dev.	Min	Pctl. 25	Median	Pctl. 75	Max
2019	2,052	771	1,007	150	414	569	842	31,936
2020	2,159	949	892	83	578	742	1,034	27,523
2021	5,592	10,542	10,155	521	4,474	8,201	14,798	405,406
2022	4,481	9,673	6,962	602	3,417	9,925	14,028	82,718
2023	1,571	1,433	1,180	112	772	1,100	1,767	23,244
<i>Total</i>	<i>15,855</i>	<i>6,821</i>	<i>8,365</i>	<i>83</i>	<i>981</i>	<i>3,822</i>	<i>11,842</i>	<i>405,406</i>


Note: The numbers are reported in USD. Source(s): Authors’ own work based on data from OpenSea.

4.5.1.2 Realized returns per parcel

Understanding individual parcel performance is crucial for investors, as it directly determines the financial outcome of market participation. Therefore, before aggregating the individual parcel returns into a market index, we examine the average realized return per parcel annually. We illustrate the geographical distribution of positive (green) and negative (red) returns per parcel on the map in Figure 4.6 and 4.7 (left-hand side) and present the annual distribution of returns using a histogram (Figure 4.6 and 4.7, right-hand side). Since our analysis does not reveal any clear locational effects on where parcels achieve the highest returns, other factors, such as market timing and sentiment may play a more significant role in explaining returns. The figures indicate that the likelihood of positive or negative returns largely depends on the year observed. For instance, the proportion of positive returns ranged from 41% in 2023 to 92% in 2021, as detailed in Table 4.2. The occurrence of heavy to total losses was notably higher in 2022 and 2023 compared to 2021. Conversely, extreme positive returns were most frequent in 2021 and 2022, a period marked by early investors potentially capitalizing on their virtual parcels amid heightened demand for digital assets and media excitement about the metaverse and virtual land. The overall dominating right-skewed distribution and the substantial discrepancies between the mean and median returns across the years highlight significant dispersion between the returns, especially extreme positive returns. However, unsold parcels could currently represent significant losses in investors’ portfolios, as our analysis is limited to realized returns and negative returns have become increasingly prevalent in 2023.


In conclusion, from an investment perspective, our findings suggest that the location of a parcel is less significant for its return potential than market timing and sentiment. This also indicates that diversification effects within the virtual land market itself may be limited. However, to validate our observations, we recommend conducting a more detailed regression analysis, particularly to explore subtle locational and other hedonic factors.

Figure 4.6: Realized returns per parcel 2019-2021



Source(s): Authors' own work based on data from Opensea and Decentraland.

Figure 4.7: Realized returns per parcel 2022-2023



Source(s): Authors' own work based on data from Opensea and Decentraland.

Table 4.2: Realized returns per parcel – Summary statistics


Year	N	Pos.	Neg.	Mean	Std. Dev.	Min	Pctl. 25	Median	Pctl. 75	Max
2019	476	0.64	0.36	0.06	0.24	-0.33	-0.04	0.04	0.13	3.88
2020	671	0.77	0.23	0.50	4.56	-0.74	0.02	0.19	0.45	117
2021	1,811	0.92	0.08	3.16	7.05	-0.63	0.11	0.38	2.39	96
2022	1,653	0.61	0.39	2.63	28.10	-1.00	-0.11	0.07	0.52	1,025
2023	573	0.41	0.59	0.28	8.27	-0.97	-0.41	-0.08	0.15	197
<i>Total</i>	<i>5,184</i>	<i>0.67</i>	<i>0.33</i>	<i>1.33</i>	<i>9.64</i>	<i>-1.00</i>	<i>-0.41</i>	<i>0.12</i>	<i>2.39</i>	<i>1,025</i>

Note: The returns and shares of positive or negative returns are presented in decimal. Exact zero returns are excluded, since we consider them to be non-arm's-length transactions and little meaningful. Note that multiple returns per parcels can occur for different years. Source(s): Authors' own work based on data from Opensea.

4.5.1.3 Number of transactions per parcel

In the context of our repeat sales index construction and to visualize market activity and investor behavior in the virtual land market, we display the geographical distribution of the number of transactions per parcel in Figure 4.8. More than half of the observed parcels (54%) have been traded more than once and 26% have been traded thrice or more. This suggests a relatively active secondary market. Geographically, transactions are widely dispersed with only vague locational effects observed in relation to the number of sales. For example, remote parcels located in the lower left of the map (around coordinates -150, -120) exhibit higher trading activity, particularly for transactions numbering four or more, while parcels adjacent to main roads in central locations demonstrate lower turnover. This might be explained by the higher parcel prices in locations with expected higher visitor density, which narrows the pool of potential buyers and reduces trading activity. Furthermore, similar to real-world properties, virtual land in prime locations is less likely to be sold frequently, if at all.

Figure 4.8: Number of transactions per parcel (n=7,665)




Source(s): Authors' own work based on data from Opensea and Decentraland.

4.5.1.4 Indices

Building on the parcel-level transaction data, this section presents the aggregated market index for virtual land, calculated using the repeat sales method by Case and Shiller (1987) and Shiller (1991). Its calculation is based on 12,325 repeat sales pairs.²⁰ The resulting index is able to explain 85% of the virtual land price changes, which highlights the credibility for its usage as previously noted in Section 4.4.²¹ The virtual land index is displayed in juxtaposition with the REIT indices in Figure 4.9.²²

Figure 4.9: Virtual land index vs REIT indices



Source(s): Authors' own work based on data from Opensea, Nareit, and S&P Global.

The virtual land market, as captured by the index, exhibits significantly higher volatility and dramatic changes when compared to the physical real estate markets. Both markets decline as a consequence of the first Covid-19 outbreak in Q1 2020. However, the virtual land market started an extreme upward trend afterwards, which can be explained by the increased interest around metaverse assets during the pandemic, as previously analyzed by Ghosh et al. (2023). Their study reveals that travel uncertainty and general social distancing induced by the Covid-19 pandemic catalyzed the growth of metaverse assets, including Decentraland's native economy token MANA. On the other hand, commercial and retail real estate markets, represented by our constructed sub-index, show the highest losses and longest time to break-even, mirroring the severe negative effects of paralyzed physical consumer and social activity. During our sample period, two pronounced peaks in the virtual land markets can be observed: first, in Q2 2021, when the interest and excitement for the metaverse and Non-Fungible Token was sparked and major brands announced their million-dollar land acquisitions on platforms like Decentraland²³ and second, in October 2022, when Facebook announced its re-branding

²⁰We also estimate the index using Bailey et al. (1963) and receive a correlation between both indices of 0.99.

²¹As a comparison, for physical real estate markets, Case and Shiller (1987) conclude that a repeat sales index can explain 7% to 27% of the variations in house prices in four US cities.

²²The gray shaded areas illustrate periods of economic uncertainty as defined by NBER: <https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions>.

²³In the place of many, see <https://www.reuters.com/markets/currencies/virtual-real-estate-plot-sells-record-24-million-2021-11-23/>.

to *Meta* associated with billion-dollar company investments into metaverse and virtual reality technologies.²⁴ While the real estate market was slowly recovering during this time, the virtual land market index experienced a 20-fold surge in value after starting to decline. At first it declined slowly, followed by a dramatic decline, which coincides with the first interest rate hikes by the FED around Q2 2022 to counteract the rising inflation rates, resulting from the ultra-loose monetary policy during the economic turmoil of the pandemic.²⁵ The resulting risk-asset-sell-off also led to significant index losses for both virtual and physical real estate markets, since both asset valuations react sensitive to interest rate elevations. With signals of pausing interest rate hikes and probable interest rate cuts, both markets began to show signals of strong recovery and a turning point in Q3/Q4 2023.²⁶

Table 4.3: Descriptive statistics of weekly returns

(a) Virtual land performance index									
Year	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 50	Pctl. 75	Max	Skewness	Kurtosis
2019	0.64	17.70	-30.43	-12.31	-1.34	10.45	55.03	0.74	0.66
2020	3.19	24.49	-39.88	-11.55	2.35	15.52	107.08	1.45	5.21
2021	7.90	23.11	-28.07	-6.47	3.76	15.71	104.61	1.60	4.85
2022	-3.83	13.83	-32.99	-10.86	-4.46	3.88	31.42	0.07	0.22
2023	2.36	20.93	-53.48	-11.65	-0.82	16.54	49.54	0.24	0.36
<i>Total</i>	<i>2.05</i>	<i>20.57</i>	<i>-53.48</i>	<i>-10.64</i>	<i>0.26</i>	<i>12.76</i>	<i>107.08</i>	<i>1.21</i>	<i>4.23</i>
(b) FTSE Nareit All Equity REITs Index USD (only Retail/Office/Gaming)									
Year	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 50	Pctl. 75	Max	Skewness	Kurtosis
2019	0.23	1.79	-3.61	-0.75	0.38	1.48	5.10	-0.30	0.32
2020	-0.07	9.37	-30.89	-4.08	0.18	2.83	30.13	0.24	3.64
2021	0.61	2.37	-5.71	-0.93	0.89	2.13	6.29	-0.25	0.42
2022	-0.49	3.28	-7.19	-2.97	-0.63	1.29	6.94	0.08	-0.21
2023	0.07	3.57	-8.67	-2.10	-0.37	2.41	9.73	0.33	0.42
<i>Total</i>	<i>0.07</i>	<i>4.88</i>	<i>-30.89</i>	<i>-1.89</i>	<i>0.14</i>	<i>2.00</i>	<i>30.13</i>	<i>0.28</i>	<i>13.95</i>

Note: This table shows the descriptive statistics for the dataset. The returns are calculated in levels and expressed as percentages. Source(s): Authors' own work based on data from Opensea, Nareit, and S&P Global.

Table 4.3 presents descriptive statistics of the weekly returns for virtual land and retail and commercial physical real estate approximated by our REIT subset. The mean return of virtual land is higher, indicating an overall higher return for investors. Virtual land exhibits a significantly wider range of returns compared to physical real estate. As already shown on a parcel-level basis, the virtual land market experiences extreme fluctuations in returns, indicating a higher level of volatility and potentially greater opportunities for both positive and negative returns, but also higher unpredictability than physical real estate markets. Investors in virtual land need to be aware of the potential for larger swings in their investment value and should carefully assess their risk tolerance before engaging in this emerging asset class.

²⁴<https://about.fb.com/news/2021/10/facebook-company-is-now-meta/>.

²⁵<https://www.cnbc.com/2022/03/16/federal-reserve-meeting.html>.

²⁶<https://www.reuters.com/markets/us/fed-flags-end-rate-hikes-sees-lower-borrowing-costs-2024-2023-12-13/>.

4.5.1.5 Correlation

Having established the performance dynamics of the virtual land market using our repeat sales index, we now present the results of the correlation analyses between virtual land returns and traditional real estate returns. The findings are critical for investors looking to diversify their portfolios with digital real estate-like assets. In the following section, we present both the temporal and frequency dimensions of this relationship. We start with the rolling Pearson correlation to capture the time-varying correlation between the returns of virtual and physical real estate markets.

4.5.1.6 Pearson rolling correlation


We calculate the Pearson correlation using an expanding rolling window of 52 weeks.²⁷ The results are depicted in Figure 4.10 and indicate overall low correlation with an average of 0.07 for the FTSE Nareit Subset returns over the whole period. The temporal correlation is highest during the outbreak of Covid-19 in Q1 2020. In line with the graphical observation of the indices, the correlation then steadily decreases since virtual land prices surge while physical real estate markets suffer from the consequences of the pandemic. As already elaborated, retail and commercial real estate markets suffered the most during this period, since physical interactions and commercial activity were limited to a minimum and digital spaces served as an alternative, resulting in increasing virtual land market returns. The distinct low correlation between virtual land and retail and commercial real estate returns is visualized through the light red area. After the correlation turns to temporary negative levels in the Q4 2021/Q1 2022 it begin to increase when the first post-Covid-19 interest rates hikes are announced by the FED, inducing a global risk-asset-sell-off and a decrease in value of both assets. This comes due to an increasing attractiveness of fixed income assets like bonds, which have higher coupon rates. REITs and real estate, in general, need to adjust to the new elevated interest rate environment, leading to higher financing costs and higher risk-free rates, ultimately resulting in lower property valuations. On the other hand, virtual land, which is primarily not financed with debt at present, was more likely to experience a decrease in prices due to its nature as a risky asset and venture investment, which generally declines in value when interest rates increase. A similar increase of correlation happened in Q3 when an end of interest rate hikes started to be expected and both assets rose in value. However, the overall correlation between virtual land and REIT returns remains quite low, averaging 0.07 with the FTSE Nareit Subset (Retail/Office/Gaming), 0.13 with the FTSE Nareit All Equity REITs Index, and 0.14 with the S&P Global REIT Index.

4.5.1.7 Wavelet analysis

The results from the wavelet analysis not only confirm the low temporal correlation observed with the Pearson correlation but furthermore uncover additional insights on the frequency dimension of the correlation, which can be interpreted as the investment horizons. Since investors consider a particular time frame when investing, it is crucial to match the investment horizon with the rate (frequency) at which certain patterns in the data occur over time.

²⁷We also calculate the Kendall's τ correlation and receive similar results. See Appendix A.6.

Figure 4.10: Virtual land and REITs (Pearson)



Source(s): Authors' own work based on data from Opensea, Nareit, and S&P Global.


For example, short-term investors might be interested in high-frequency daily or weekly correlation movements of two assets, while long-term investor would care more about broader correlation trends that develop over many months and years. Figure 4.11 represents the wavelet analysis results in a three-dimensional coordinate system, where the z-axis displays the correlation (wavelet coherence) from 0 (lowest correlation) to 1 (highest correlation).²⁸ Its intuition can be interpreted like the squared coefficient of correlation. The color map helps to better represent the correlation and is simultaneously used in the two-dimensional plot in Figure 4.12. The horizontal axis represents time, while the vertical axis represents the investment horizons, converted to weeks. The *cone of influence* indicated by the dotted lines forming a triangle, indicate areas of less confidence and reliance in results toward the edges. This is due less data available for the calculations at the beginning and the end of the time series.

In the following, we present the bivariate wavelet analysis considering the virtual land market returns and the FTSE Nareit Subset, which focuses on retail and commercial real estate. Several interesting findings within the cone of influence can be identified. Firstly, similar to the Pearson correlation coefficients, there is a general low co-movement between virtual land market returns and REIT returns, which is indicated by predominant blue colors and valleys in Figure 4.11 and 4.12. However, we find two particularly very pronounced correlation spikes in the medium-term investment horizon during the Covid-19 outbreak in early 2020 and the first FED interest rate hikes in mid-2022. For the short-term investment horizon, we find some minor spikes of correlation associated with short-lived general stock market movements, apparently affecting virtual land and equity REITs equally. The findings for the FTSE Nareit All Equity REITs Index and the S&P Global REIT Index yield similar results and are presented in Appendix A.5.

Starting with the highest frequency bands of the wavelet analysis (around 0–10 weeks), increased co-movement between both markets can be observed, suggesting that since both


²⁸We performed the wavelet coherence analysis for the other REIT indices as well and received similar results. Please refer to Appendix A.5.

Figure 4.11: Virtual land & FTSE Nareit Subset



Source(s): Authors' own work based on data from Opensea and Nareit.

Figure 4.12: Virtual land & FTSE Nareit Subset



Source(s): Authors' own work based on data from Opensea and Nareit.

markets are frequently traded and risk-oriented, they are significantly influenced by media, short-term news, and trading activities. Nevertheless, these periods of high co-movement appear to be short-lived. Additionally, previous studies have shown that REITs tend to comove more with equity markets than with private real estate in the short term, resulting in higher volatility (see, e.g., Ling and Naranjo, 1999).

Moving to the frequency bands around 10–30 weeks of the wavelet analysis, which is considered the mid-term investment horizon, a distinct region of very high co-movement in the time-frequency space can be observed during the Covid-19 outbreak in early 2020. During this time and frequency, the co-movement between both market returns was nearly 1, as indicated by the dark red color. This has several implications. Firstly, extreme events like the Covid-19 outbreak tend to result in high correlation between REITs and the virtual land market observed, reducing the safe haven benefits of both assets.²⁹ However, there does not seem to be a long-lasting regime switch afterwards. On the contrary, we observe that during the period of Covid-19 recovery, the correlation between both assets reaches its lowest values, underpinning our previous findings: during lockdown and social distance, the demand for metaverse technology and digital assets, such as virtual land increases while commercial and retail properties suffer. The correlation remains low, until the FED announces interest rate hikes in mid-2022, indicating that both assets equally are influenced by major global macroeconomic events.

For the lowest frequency bands above 30 weeks, which is considered the long-term investment horizons, we observe a generally low correlation within the cone of influence, suggesting potential hedging properties of virtual land.³⁰ But it also becomes observable that the correlation began to increase in mid-2022, when the FED ended the low interest rate environment.³¹ Even though this observation lies outside the cone of influence and may not be entirely reliable, it suggests that both assets tend to increasingly co-move with the macroeconomic environment, probably underpinning its increasingly legitimacy as a complementary investment class. This could indicate the potential for diminishing hedging effects in the future, though these changes are not yet fully significant.

Summarizing, our correlation analyses provide evidence, that virtual land assets could be of interest to real estate investors seeking low-correlated retail or commercial-like properties in the digital realm, particularly those with a higher risk tolerance. On average, virtual land exhibits low correlation with traditional real estate assets. However, our results show that there are acute periods of spiking correlation during market turmoil, such as the immediate Covid-19 outbreak, interest rate regime shifts and occasional short-term market reactions. When looking for long-term diversification, investors should pay attention to an increasing co-movement between both assets in the future, as virtual land may mature and become more integrated with the broader economy.

²⁹Following Baur and McDermott (2010) and Ko and Lee (2023), we define a safe haven as an asset that shows a low or negative correlation with another asset during periods of economic turmoil.

³⁰Based on the definitions by Baur and McDermott (2010) and Ko and Lee (2023), we define a hedge as an asset that exhibits a low or negative correlation with another asset on average over the long term.

³¹Note that although we are discussing “long-term” investment horizon here, we are aware that in traditional real estate, this usually spans more than a year.

4.5.1.8 Regression analysis

Having examined the correlation structure between virtual and physical real estate, we now focus on determining the factors that drive and explain it.³² By understanding these factors, investors can more accurately predict, how virtual land might behave in relation to traditional real estate under varying market conditions and adjust their portfolios and investment decisions accordingly. To gain insights into potentially significant variables, we adopt the methodology of Christoffersen et al. (2012) by conducting regression analyses on a broad selection of variables. Given the novelty of this paper’s investigation to estimate and explain these correlations, there is a lack of existing literature on potential determinants. However, as a starting point, we consider variables from related studies on real estate token (see Kreppmeier et al., 2023) and cryptocurrencies (see, e.g., Mou et al., 2024) and then extend and modify them to our setting. Our variables encompass macroeconomic determinants, crypto-related information, and sentiment-based indices.³³ To facilitate the selection process, we follow the standard approach of step-wise regression, where variables are included in the model only if they improve the selection criteria. In this analysis, we employ the Akaike Information Criteria (AIC) as a commonly used criterion. Given that we observe correlations along both the frequency and time dimensions, $R^2(\tau, s)$, we incorporate quarter-year time fixed-effects and treat frequency as a trend variable in each regression, following Christoffersen et al. (2012). Additionally, we address potential issues of heteroskedasticity and autocorrelation by employing Heteroskedasticity and Autocorrelation Consistent (HAC) Standard Errors. To determine the appropriate number of lags for this purpose, we refer to Newey and West (1987) and select four lags in our case. To allow for a direct comparison of the variables’ impacts, we standardize all of them within each regression analysis.³⁴ To gain a more detailed understanding of the determinants, we partition the data along the frequency dimension into three segments: a short-term investment horizon spanning 0–10 weeks, a medium-term investment horizon from 10 to 30 weeks, and a long-term investment horizon exceeding 30 weeks. This segmentation enables us to identify the driving factors behind the connectedness across various frequencies. We present the final regression estimates for each sample, as well as the entire pooled sample, in Table 4.4.

Starting with the frequency trend, we observe a significant negative estimate for short- and medium-term correlations and the overall time frame, suggesting an overall increase in diversification effects over time. Moving on to the short-term horizon, we identify several macroeconomic variables as significant determinants. Firstly, we find a significant positive relationship with the Consumer Sentiment Index and the Aruoba-Diebold-Scotti Business

³²We also performed regression analyses to explain the returns of the virtual land and REIT indices, which yielded fewer significant relationships than those found in the correlation structure analysis. This finding is consistent with the expectation that returns are inherently more difficult to predict due to their higher volatility and unique, asset-specific factors that are often unpredictable. In contrast, the correlation structure tends to be more stable over time as it captures the influence of broader macroeconomic conditions or market sentiment that affect both assets simultaneously, which makes it relatively easier to identify its underlying drivers of this correlation. Particularly in the case of explaining virtual land returns, future research is encouraged to delve deeper into the return structure of this emerging market to better understand its complexity and dynamics.

³³For clarity, Table A.6 provides a list of the variables used, accompanied by brief descriptions and their respective data sources.

³⁴In this context, standardizing refers to converting each variable to a Z-score by subtracting the mean and dividing by the standard deviation. This process ensures that all variables are on a common scale, facilitating a more straightforward comparison of their relative effects.

Table 4.4: Determinants of correlation estimates

	Frequency horizon			All data
	Short frequency (0–10 weeks)	Medium frequency (10–30 weeks)	Long frequency (>30 weeks)	
Intercept	0.284***	0.309***	0.407***	0.332***
Frequency trend	-0.030***	-0.067***	0.088***	-0.028***
<i>Macroeconomic data</i>				
Consumer Sentiment Index	0.099***	-0.033***	0.014	0.031**
US Economic Uncertainty Index	0.020***	0.016**		0.009*
Global Economic Uncertainty Index	0.017**	0.009		0.010**
World Uncertainty Index	-0.010	0.011**	-0.007	
ADS Index	0.018***			0.006
Δ Fed Funds Rate		0.003		
Term spread	-0.015	-0.018		-0.012
<i>Crypto data</i>				
MANA (USD)	-0.057***			-0.020**
ETH (USD)	0.019	-0.011	0.017*	
Volatility of MANA	0.018***		-0.006*	0.005**
Volatility of ETH	-0.012*	-0.009*		
ETH gas fees (USD)		-0.007**		
<i>Sentiment data</i>				
Google Search “Metaverse Real Estate”	0.021***			0.006
Google Search “Buy Real Estate”	-0.041***		-0.005	-0.017***
R^2	0.165	0.736	0.480	0.111
Adjusted R^2	0.161	0.735	0.478	0.110
AIC	-3,434	-6,689	-6,596	-7,248
Observations	6,264	4,698	5,220	16,182

Note: This table presents the regression estimates for the variables determining the correlation between virtual land market and physical real estate market returns. The data are segmented into short-term (0–10 weeks), medium-term (10–30 weeks), long-term (>30 weeks), and the entire sample period regardless of the frequency (All data). The regression uses standardized variables, quarter-year fixed effects, and heteroskedasticity- and autocorrelation-consistent (HAC) standard errors with four lags. The model selection is based on the Akaike information criterion (AIC). For the description of variables and detailed data sources please refer to Table A.6 in the Appendix. Source(s): Authors’ own work.

Conditions Index (ADS Index), implying that the correlation increases during prosperous times and decreases during economic downturns. In other words, when the economy flourishes (deteriorates) the interest of investors for both physical and virtual real estate increases (decreases). The variables of US and global economic uncertainty support these findings that greater economic uncertainty leads to a higher correlation. In other words, in the short-term the co-movement of both assets tends to be connected to the macroeconomy equally sided.

Moving on to crypto-related determinants, we find a significant negative impact of MANA, Decentraland's economy cryptocurrency. This can be explained by the fact that as the price of MANA increases, the investment in virtual land becomes more expensive and less attractive for investors. Our observation that 72 % of the virtual land is bought with MANA supports this. Overall, this results in a decrease in correlation. Similarly, the significant positive sign of MANA's volatility underpins this trend. When the volatility of Decentraland's cryptocurrency increases, investors may be inclined to reduce their investment in virtual land. Next, we examine the estimated signs of the Google Search Volume. We observe a significant positive sign for the search term "Metaverse Real Estate" and significant negative sign for "Buy Real Estate". This suggests that the correlation between virtual and physical real estate increases when investors' search volume focuses on virtual land in the metaverse. This co-movement might indicate that when investors are interested in metaverse real estate, they are simultaneously engaged with physical real estate markets, potentially viewing virtual land as an extension of their investment activities into the digital realm. Conversely, when search interest shifts predominantly toward physical real estate, as indicated by increased searches for "Buy Real Estate", the correlation decreases, reflecting a demand shift toward tangible real estate over the virtual one. Overall, the Consumer Sentiment Index exhibits the largest (absolute, standardized) beta coefficient, implying that the perception of consumers regarding their financial situation and the general state of the economy plays the most significant role as a determinant of correlation in the short-term investment horizon. The second most important determinant is the MANA price, which is Decentraland's economic currency token. Consequently, the prevailing economic climate and the cost of investment and appear to be the dominant drivers of the correlation between virtual and physical land in short-term horizons.

For medium-term investment horizons, we observe opposite estimated signs for the Consumer Sentiment Index. Notably, the correlation between both market returns is highest in the medium-term investment horizons during Covid-19, which was characterized by low consumer sentiment due to lockdown and social containment. On the other hand consumption of digital products in the virtual world flourished during this phase. In other words, while the investment in retail and commercial real estate faltered, digital products like virtual land gained attraction among investors and the correlation between both assets decreased due to one-sided demand shifts. However, as stated out before, the immediate outbreak of Covid-19 indicated by an increased World Uncertainty Index and Global Economic Uncertainty Index resulted in a prior increase in the correlation between both market returns.

Among the crypto-related data, we find the ETH Gas fee to be a significant determinant, with a negative sign. This implies that as transactions on the Ethereum blockchain become more expensive, the correlation decreases. Higher costs make it less favorable for investors to invest in virtual land, leading to a decrease in correlation. Overall, the medium-term correlation is primarily driven by the consumer sentiment, as indicated by the large beta coefficients, and less influenced by crypto or attention-related determinants.

For long-term horizons, we do not find any of the economic or attention-related determinants to be significant. In contrast to short- or medium-term correlations, the long-term correlations seem unaffected by different economic climates or the prevailing attention surrounding both asset classes. Initially, this might seem disappointing, but upon closer examination, we can derive a positive inference for investors: the correlation remains robust at a consistently low level across various market phases. Interestingly, the ETH price shows a slight positive effect on long-term correlation. This increase in value may indicate a broader acceptance and adoption of blockchain technology. Given that Decentraland and most virtual land platforms are built on the Ethereum platform, virtual land as an investment class could benefit from this, potentially enhancing its correlation with traditional real estate markets. Furthermore, in the long-term horizon, the frequency trend as an explanatory variable is significantly positive. This supports our observation that over the long term, the returns on virtual and physical real estate tend to correlate more, suggesting that virtual land is gaining recognition as a viable new asset class among the broad economy.

If we conduct our analysis on the entire sample, we find that the consumer sentiment emerges as the most crucial determinant, underscoring the impact of consumption behavior on retail and commercial real estate as well as virtual real estate. This seems logical, as both assets serve as platforms that fulfill consumer needs. The frequency trend emerges as the second most crucial determinant. For the overall sample, it indicates enhanced diversification advantages for generally longer investment horizons. Additionally, the MANA price and the attention given to physical real estate investment are important determinants that can explain the correlation between both asset types. Surprisingly, changes in the Fed Funds Rate cannot statistically explain the correlation.

4.5.1.9 Portfolio implications

This final section builds on the previous findings on correlation patterns and examines the practical implications of adding virtual land to an investor's portfolio. By analyzing various portfolio strategies, both with and without the inclusion of virtual land, we aim to provide insights into how this emerging digital asset class could impact overall portfolio performance. We construct various portfolios based on mean-variance optimization (see Markowitz, 1952), 1/N naive portfolio diversification (see DeMiguel et al., 2009) and economically reasonable portfolio allocations that include a 10%–20% portion of virtual land into a traditional real estate or multi asset portfolio.³⁵ To compare portfolio performance, we examine a set of metrics, widely used in portfolio management: expected return (μ), standard deviation (SD), and Sharpe ratio $\frac{\mu - r_f}{SD}$, where r_f is the return of the risk-free asset (Sharpe, 1994). To more precisely account for the high volatility of virtual land returns, we incorporate the Sortino ratio (Sortino and Meer, 1991).³⁶ The Sortino ratio is calculated as $\frac{\mu - r_f}{SD_{downside}}$.³⁷ The results of the portfolio metrics for the respective portfolios are presented in Table 4.5. We compare

³⁵We acknowledge that allocating 10%–20% of virtual land in an investor's portfolio may still be considered as economically unreasonable. However, this allocation is presented as part of an initial empirical investigation, based on the assumption that the investor considers virtual land a viable asset class.

³⁶Unlike the Sharpe ratio, the Sortino ratio adjusts for the assumption of normally distributed returns by penalizing only downside deviations. This adjustment is particularly relevant given that virtual land returns are positively skewed and often include extreme positive returns, which should not be penalized as they contribute favorably to the overall performance.

³⁷In our calculations we set the minimum acceptable return (MAR) to r_f .

the portfolio strategies by examining each strategy both with the inclusion of virtual land (+Virtual land) and without it (Benchmark). We find that the addition of virtual land consistently improves the expected and risk-adjusted returns of portfolios, as measured by the Sharpe and Sortino ratios. For example, in the REITs portfolio, the inclusion of 10% virtual land increased the Sharpe ratio from 0.10 to 0.44 and the Sortino ratio from 0.09 to 0.44. However, due to its inherent volatility, the inclusion of virtual land led to higher portfolio standard deviations. Interestingly, in the REITs and Commercial REITs portfolios, adding 10% virtual land resulted in only a marginal increase in standard deviations (from 0.26 to 0.29 and from 0.35 to 0.36), while significantly boosting the expected returns from 4% to 15% and from 3% to 14%, respectively. Therefore, adding virtual land to a traditional real estate portfolio represented by REITs or Commercial REITs can unfold the greatest relative benefits in terms of risk-adjusted portfolio returns.³⁸

³⁸Further details on the performance metrics by year and asset and the correlation matrix of excess returns of the observed assets can be found in Table A.7, A.8, and A.9 in the Appendix.

Table 4.5: Comparison of portfolio performance metrics (annualized)

Strategy	Portfolio	μ	SD	Sharpe	Sortino	Weight Virtual land
1/N	Benchmark	0.05	0.20	0.15	0.15	0.20
	(+Virtual land)	<i>0.25</i>	0.36	<i>0.66</i>	<i>0.76</i>	
Max Sharpe	Benchmark	0.12	0.19	0.53	0.50	0.17
	(+Virtual land)	<i>0.28</i>	0.33	<i>0.80</i>	<i>0.94</i>	
Max Sortino	Benchmark	0.12	0.19	0.53	0.50	0.40
	(+Virtual land)	<i>0.50</i>	0.62	<i>0.76</i>	<i>0.99</i>	
Min variance	Benchmark	0.04	0.12	0.19	0.19	0.00
	(+Virtual land)	0.04	0.12	0.19	0.19	
REITs	Benchmark	0.04	0.26	0.10	0.09	0.10
	(+ 10% Virtual land)	<i>0.15</i>	0.29	<i>0.44</i>	<i>0.44</i>	
	(+ 20% Virtual land)	<i>0.25</i>	0.39	<i>0.60</i>	<i>0.66</i>	
Commercial REITs	Benchmark	0.03	0.35	0.05	0.05	0.10
	(+ 10% Virtual land)	<i>0.14</i>	0.36	<i>0.34</i>	<i>0.34</i>	
	(+ 20% Virtual land)	<i>0.24</i>	0.42	<i>0.53</i>	<i>0.57</i>	
Multi asset	Benchmark	0.09	0.16	0.42	0.39	0.10
	(+ 10% Virtual land)	<i>0.18</i>	0.23	<i>0.72</i>	<i>0.78</i>	

Note: This table compares the annualized performance measures: mean return (μ), standard deviation (SD), Sharpe ratio and Sortino ratio for various portfolios with and without the inclusion of virtual land over the sample period from January 2019 to December 2023. The REITs portfolio is represented by the S&P Global REIT Index (USD). The Commercial REITs portfolio is represented by our constructed FTSE Nareit All Equity REITs Index USD (only Retail/Office/Gaming) (see Section 4.3.2). The multi-asset portfolio consists of 70% stocks (MSCI World Index USD), 20% bonds (FTSE World Broad Investment-Grade Bond Index USD), and 10% REITs (S&P Global REIT Index USD). When virtual land is included, the stock exposure reduces to 60%. All portfolio weights for each asset in each portfolio strategy are presented in Table A.10 in the Appendix. The expected returns and the standard deviations are calculated using weekly returns. The ratios are calculated using weekly excess returns over the risk-free rate (r_f), which is the 3-Month Treasury Bill Secondary Market Rate (<https://fred.stlouisfed.org/series/DTB3#0>). Italic numbers highlight the performance improvement for the respective metric when virtual land is included in the portfolio. Source(s): Authors' own work based on data from Opensea, Nareit, S&P Global, FTSE, MSCI, and FRED.

4.6 Conclusion

In this study, we analyzed the virtual land market of a pioneering and popular blockchain-based virtual world that has attracted numerous global companies and investors to purchase purely digital land and explore new business potentials. Given the intrinsic parallels between virtual and physical real estate markets, our main objective was to explore the relationship between the returns of both markets and provide practical implications for market participants and investors. We compiled a unique dataset of virtual land sales between January 2019 and December 2023 and constructed a virtual land market index using repeat sales index methodology from traditional real estate studies. Based on the resulting return series for virtual land, and various equity REIT indices, we employed wavelet coherence analysis to examine their dynamic correlation across time and different investment horizons. To further understand the correlation patterns, we performed regression analyses on various macroeconomic, crypto-related, and sentiment-based variables. Lastly, we analyzed different portfolio strategies with the inclusion of virtual land and compared their performances.

Our results reveal that the correlation between virtual and physical land returns is generally low, with the lowest correlation observed in the medium term investment horizon between two to eight months and during the Covid-19-induced lockdown period starting in mid-2020, when physical consumer activity, travel and social interactions were paralyzed. The correlation is highest in short-term extreme events, such as the acute Covid-19 outbreak in 2020, interest rate change announcements, and short-term stock market volatility up to two months, which is often observed for equity REITs. During these periods of significant turmoil and overall market uncertainty portfolio diversification effects and safe haven properties vanish, and both markets exhibit significant co-movement. We find that the correlation is primarily driven by the consumer and economic climate, the price of the virtual economy token and the investors' attention regarding both assets. Our portfolio analysis indicates that virtual land can enhance risk-adjusted returns within a traditional portfolio, especially when added to a commercial real estate portfolio. However, it is important to note that virtual land remains a nascent and highly volatile niche asset class.

While our study provides valuable initial insights into virtual land markets and its performance, it has several limitations. The most significant limitation is the relatively short time frame available. With extended data, it would be possible to broaden our analysis to include private real estate markets to reinforce our findings, given that REITs are an indirect proxy of real estate markets. Another limitation is the approximation of virtual land markets using only one metaverse platform (Decentraland). Although we argue that it provides the closest representation of its physical counterpart and offers the richest data set, incorporating multiple virtual land platforms would enhance the robustness of our results.

Future research could include a comparative analysis between different virtual land markets and reveal crucial distinctions between their functionalities. Additionally, a more detailed investigation of institutional investor attitudes toward virtual land markets can provide important insights into their maturity and future trajectory. Although still in its early stage, virtual land markets could become an integral part of the growing intersection of blockchain and the real estate ecosystem. Therefore, the impact of virtual land markets on physical real estate, particularly in terms of potential capital shifts and their ability to create parallel and more efficient transaction layers for physical real estate through blockchain technology, holds

significant research potential. Likewise, the broader societal impact of virtual land markets has yet to be thoroughly examined. As more data and platforms become available, future studies are encouraged to use our initial empirical findings as a foundation. Thanks to public blockchain technology, executing new research is now more accessible than ever.

Appendix

A.1 Global companies in Decentraland

Table A.1: Global companies and XY-coordinates in Decentraland (as of 10 April 2024)

Company (Country)	Business Sector	X	Y
19th Hole (US)	sports & entertainment	-130	78
Adidas (Germany)	fashion & clothing	-87	-2
Aktif Bank (Turkey)	banking & financial services	64	-126
Amorepacific (South Korea)	cosmetics & beauty	-109	-5
Arab National Bank (Saudi Arabia)	banking & financial services	-35	128
ArentFox Schiff (US)	law & consulting	-100	-21
Auroboros (UK)	fashion & clothing	-95	-13
Australian Open (Australia)	sports & entertainment	-127	71
Avalanche	crypto	12	-16
Balenciaga (France)	fashion & clothing	-79	-9
Banco de Chile (Chile)	banking & financial services	48	-41
Ben Bridge (US)	fashion & clothing	-108	-17
Binance.US (US)	crypto	71	23
Bitcoin Suisse (Switzerland)	banking & financial services	54	26
Bojangles (US)	food & beverage	-108	-13
Campbell Company of Canada (Canada)	food & beverage	104	-141
Coach (US)	fashion & clothing	-76	-7
Commercial Bank International (UAE)	banking & financial services	63	-128
Consorsbank (Germany)	banking & financial services	57	28
dating.com (Malta)	other	127	-121
Deka Bank (Germany)	banking & financial services	-5	-128
Deloitte (US)	law & consulting	-19	-132
Deutsche Bank (Germany)	banking & financial services	-63	1
DKNY (US)	fashion & clothing	-106	-18
Dolce & Gabbana (Italy)	fashion & clothing	-95	-19
Domino's Pizza (US)	food & beverage	-92	104
Doritos by PepsiCo (US)	food & beverage	-67	-119
Dundas (US)	fashion & clothing	-95	-15
Emirates (UAE)	other	-131	67
Energias de Portugal (Portugal)	other	-71	-121
EY (UK)	law & consulting	58	127
Fidelity (US)	banking & financial services	-18	144
FileCoin	crypto	-5	-132
French's Army (France)	universities & institutions	-18	-110
Gleis Lutz (Germany)	law & consulting	42	-55
Heineken (Netherlands)	food & beverage	-122	-82
Hong Kong University (Hong Kong)	universities & institutions	-19	102
KB Home (US)	real estate & construction	-111	-19
Kraken (US)	crypto	91	13
Kraken (US)	crypto	-62	75
Kraken (US)	crypto	-91	-7
La Liga (Spain)	sports & entertainment	-149	64
Lancia (Italy)	other	-102	-12
Love Honey (UK)	cosmetics & beauty	-49	-106
Lush (US)	cosmetics & beauty	-56	-110
Majid Al Futtaim (UAE)	real estate & construction	-30	-28

Continued on next page

Table A.1 continued from previous page

Company (Country)	Business Sector	X	Y
Maker	crypto	68	25
Maybelline New York (US)	cosmetics & beauty	70	-146
Miami Fashion Week (US)	fashion & clothing	-11	-49
Moiq Capital (Singapore)	banking & financial services	-8	-45
Monnier Paris (France)	fashion & clothing	-95	-17
Mountain Dew by PepsiCo (US)	food & beverage	-71	-124
MTS (Russia)	telecommunication	66	-148
Naked Life (Australia)	food & beverage	58	25
National Payments Corporation of India (India)	banking & financial services	-130	-134
Nestlé (US)	food & beverage	-57	-57
Neutrogena (US)	cosmetics & beauty	-101	88
Next Realty (Brazil)	real estate & construction	-75	35
NTT Data (Japan)	law & consulting	41	-47
Oversea-Chinese Banking Corporation (Singapore)	banking & financial services	-20	-111
Polygon	crypto	55	19
PwC (US)	law & consulting	-127	-115
Quontic Bank (US)	banking & financial services	-30	30
Riyad Bank (Saudi Arabia)	banking & financial services	61	-124
Sotheby's (US)	other	52	83
Swiss Prime Site (Switzerland)	real estate & construction	55	125
Sygnium Bank (Switzerland)	banking & financial services	-8	12
The Voice Studios (US)	sports & entertainment	-97	100
Tokens.com (US)	crypto	69	30
Toledo Museum of Art (Spain)	universities & institutions	48	-49
Tommy Hilfiger (US)	fashion & clothing	-107	-15
Touch (Lebanon)	telecommunication	66	29
Transak ATM (US)	crypto	57	29
Umdasch Group (Netherlands)	real estate & construction	-34	-4
Universidad Itaca (Mexico)	universities & institutions	-67	1
VfL Wolfsburg (Germany)	sports & entertainment	110	-26
William Mattar Law Offices (US)	law & consulting	-24	-5
Worldline (France)	banking & financial services	52	27

Source(s): Authors' own work based on data from Decentraland.

Figure A.1: Companies in Decentraland - Exemplary excerpt (as of 10 April 2024)



Note: Companies in Decentraland from left to right: Adidas, Aktif Bank, Arab National Bank, Hong Kong University, ArentFox Schiff, Australian Open, Banco de Chile, National Payments Corporation of India, Commercial Bank International, Dating.com, Coach, NTT Data, Swiss Prime Site, Majid Al Futtaim, Tommy Hilfiger, Quontic Bank, Deka Bank, EY, Mountain Dew by PepsiCo, Riyadh Bank, Consorsbank, Dolce & Gabbana, Heineken, Sotheby's, Deutsche Bank, FileCoin, Oversea-Chinese Banking Corporation, Toledo Museum of Art, Energias de Portugal, La Liga, Doritos (by PepsiCo), Tokens.com, Binance.US, Kraken, Deloitte, Sygnum Bank, Touch, Miami Fashion Week, PwC, Gleis Lutz. Source(s): Authors' own work based on data from Decentraland.

Table A.2: Number of companies by business sectors

Business sector	Number of companies
banking & financial services	16
fashion & clothing	11
food & beverage	8
consulting	7
crypto	6
universities & institutions	6
sports & entertainment	5
cosmetics & beauty	5
other	5
real estate & construction	5
telecommunication	2
<i>Total</i>	<i>76</i>

Source(s): Authors' own work based on data from Decentraland.

Table A.3: Number of companies by continent

Continent	Number of companies
North America	30
Europe	25
Asia Pacific	8
Middle East	6
Other	4
South America	3
<i>Total</i>	<i>76</i>


Source(s): Authors' own work based on data from Decentraland.

A.2 Virtual land market – Further analysis

A.2.1 Transaction activity

In this section, we broaden the virtual land market analysis by examining the transaction volume over time and the investor participation. The monthly transaction volume and the number of transactions for virtual land parcels are displayed in Figure A.2, Table A.4 and Table A.5. The total dollar transaction volume for all virtual land parcel transactions amounts to \$199,198,077. Fueled by inflated expectations surrounding the metaverse, extensive media coverage, and a possible increase in online activity due to Covid-19, 2021 and 2022 saw a dramatic surge in transaction activity, with average monthly volumes peaking at \$8,914,267 in 2021 and \$6,844,938 in 2022. Although this elevated activity proved unsustainable and the transaction volumes sharply decline in 2022, the volume in 2023 remained above pre-hype levels, averaging \$353,623 per month compared to \$220,899 in 2019 and \$266,110 in 2020. Notably, although the dollar transaction volume declined significantly in 2022, the number of transactions decreased less dramatically, with the average monthly transactions in 2023 stabilizing at 151, compared to the all-time high average of 547 in 2021. Although at significantly lower price levels, this indicates that the market remains moderately active. Continued virtual land purchases by global companies in 2023 underscore the sustained interest of virtual land platforms and the metaverse as a whole.

Figure A.2: Monthly transaction activity



Source(s): Authors' own work based on data from Opensea.

Table A.4: Average monthly dollar transaction volume – Summary statistics

Year	Mean	Std. Dev.	Min	Pctl. 25	Median	Pctl. 75	Max
2019	220,899	107,206	100,734	142,407	188,720	284,537	418,555
2020	266,110	200,303	148,762	177,062	217,801	241,349	885,782
2021	8,914,267	11,773,831	731,574	3,041,820	4,354,482	6,900,455	33,883,825
2022	6,844,938	9,100,382	707,819	979,361	2,132,500	8,132,239	28,810,880
2023	353,623	284,012	144,571	194,000	270,085	400,791	1,182,611
Total	3,319,967	7,472,785	100,734	202,366	392,039	2,086,066	33,883,825

Note: The numbers are presented in USD. Source(s): Authors' own work based on data from Opensea.

Table A.5: Average monthly number of transactions – Summary statistics

Year	Mean	Std. Dev.	Min	Pctl. 25	Median	Pctl. 75	Max
2019	193	76	109	141	172	207	343
2020	203	121	117	149	175	189	567
2021	547	370	194	299	436	592	1,289
2022	424	287	174	222	317	545	1,141
2023	151	38	97	130	150	154	249
Total	304	263	97	152	187	347	1,289

Source(s): Authors' own work based on data from Opensea.


A.2.2 Investor participation

To provide further insights into the structure of the virtual land market, we present a breakdown of various investors categories based on their average dollar transaction volume (Figure A.3). We categorized 9,949 unique wallet addresses (traders) from our transaction data into four distinct investor groups using psychological dollar thresholds. The largest group (48.1%) consists of investors trading virtual land between \$1,000 and

\$10,000 on average, referred to as *Crypto Savvy Investors*. We assume this group to have a solid understanding of the virtual land market and to be willing to allocate a substantial portion of their portfolios into it. Another significant group, making up 36.9% of all investors, consists of *Expert Investors*, whose average transaction amounts are between \$10,000 and \$100,000. We classify these investors as experts in the virtual land market due to their higher level of investment. Investors with average transaction amounts below \$1,000 are classified as *Retail/Small Investors* and make up 14.0% of all investors. This group likely includes individuals who are curious about virtual land but invest smaller amounts, reflecting their limited experience in this market. Finally, investors transacting more than \$100,000 on average are classified as *Institutional Investors*. These investors, though comprising only 1.1% of the total, are likely corporate entities with the capital and confidence to invest such large sums. Additionally, they may include developers of virtual land platforms who were able to acquire land during the earliest stages of market development.

In summary, the observed virtual land market appears to be dominated by experienced investors who are knowledgeable about this emerging asset class and willing to invest significant amounts. However, due to the price volatility the boundaries we set for the investor groups are subject to changes. While further research into investor participation and the structure of virtual land markets goes beyond the scope of this paper, we encourage it for future research.

Figure A.3: Investor groups by average dollar transaction size



Source(s): Authors' own work based on data from Opensea.

A.2.3 Market development and further implications

The Hype Cycle Model by Gartner (Fenn and Raskino, 2007) provides a useful framework for assessing the long-term trajectory of virtual land markets as an innovative platform moving towards mainstream adoption. The model outlines the evolution of technological innovations through key phases pronounced by the peak of inflated expectations, disillusionment, recovery of expectations, and finally real-world productivity. Based on the model, the current state of virtual land markets appears to be in the phase of disillusionment, where slower-than-expected adoption and missed expectations dominate. For example, Decentraland, with approximately 7,000 monthly active users, has fallen significantly short of market projections. However, the growing number of corporate initiatives on virtual land platforms in the metaverse, accelerating technological advancements in virtual reality and blockchain, and the growing equivalence of virtual and physical interactions, particularly among younger audiences (see Westcott et al., 2023), set the path for a recovered and growing market in the coming years. If even a tiny fraction of the trillion-dollar physical real estate market is captured by the virtual real estate markets as an alternative platform to fulfill customer and social needs, the market potential is enormous. By 2030 alone, the market for virtual real estate is expected to grow from the current \$2.5 billion to \$8 billion.³⁹ Achieving this growth, however, will require addressing key challenges. The infrastructure for

³⁹See <https://www.statista.com/outlook/amo/metaverse/metaverse-virtual-assets/worldwide>.

creating decentralized (Web3) virtual worlds and virtual land markets must be enhanced, particularly in terms of their graphical quality and immersiveness. Additionally, the lack of standardization and interoperability between virtual platforms, where digital blockchain-based assets can be transferred across platforms, remains a significant barrier. Only with these problems solved and clear regulatory frameworks, institutional capital market participation and mainstream adoption could be fully realized. However, when for example considering the recent approval of a spot Ethereum ETF by the US Securities and Exchange Commission (SEC), which is the underlying blockchain of Decentraland, we believe that an institutionalization of other emerging assets like virtual land is possible in the future.⁴⁰

A.3 Repeat sales index

Traditional repeat sales index methods (see, e.g., Bailey et al., 1963; Case and Shiller, 1987) are designed for real world properties that sell twice (resale) in the observed samples and rarely more frequent. Half of our observed resold virtual parcels are traded thrice or more. To avoid overlapping sale intervals, we follow the proposed procedure by Shiller (1991) and treat multiple resales as independent observations. He proposes that in the case of three sales of a property, the sales pairs are constructed (first, second) and (second, third). Thus, with n transactions of one parcel, it results $(n - 1)$ in repeat sales pairs.

We start elucidating the repeat sales index construction by writing the basic repeat sales model by Bailey et al. (1963) in matrix notation with

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{U}. \quad (4.2)$$

$\mathbf{Y}$ is the dependent variable vector, containing the log price changes of our sales pairs. $\mathbf{X}$ is the independent time dummy variable matrix filled with 0 (no sale), -1 (first sale) and 1 (second sale). $\mathbf{U}$ is the vector of errors. The errors are assumed to have a log-normal distribution; i.e., that they have zero means and the same variances σ^2 . Thus, the log error term can be expressed as follows: $\log \mathbf{U} \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$, whereby *iid* means independent and identically distributed. The error variance is assumed to be constant or homoscedastic and thus not dependent on the length of time between sales as in the method by Case and Shiller (1987). By running an OLS regression, one can estimate the coefficient vector $\hat{\boldsymbol{\beta}}$.

We account for the number of repeat sales for each parcel by adding weights to the regression and therefore split the effect of each sales pair with a sale that appears in another sales pair across sales pairs. For instance, the weight (ω_i) for a sales pair, which sales appear in three sales pairs, would then be 0.33. We incorporate the weighting vector $\boldsymbol{\Omega}$ into the regression and obtain the following index construction methodology denoted in matrix notation

$$\mathbf{Y} = \boldsymbol{\Omega}\mathbf{X}\boldsymbol{\beta} + \mathbf{U}. \quad (4.3)$$

The repeat sales method by Case and Shiller (1987) is a generalized-least-square regression (GLS); i.e., compared to the OLS regression, initially proposed by Bailey et al. (1963), it assumes unstable error variance across houses (heteroscedastic errors). So, the holding period between two sales of the same property should increase the variance of the prices due, for example, to random changes in neighborhood quality or random differences in the amount of refurbishments across houses.⁴¹ The log error term to be the sum of the individual contributions from each sale ($\sigma_u^2 + \sigma_u^2$) plus a gaussian random walk representing the time periods between sales t and s , which is $((t - s)\sigma_v^2)$. It results in the following log error term $\log \mathbf{U} \sim \mathcal{N}(0, 2\sigma_u^2 + (t - s)\sigma_v^2)$.

Case and Shiller (1987) propose the following three-stage procedure:

1. A price index is estimated, as has been proposed originally by Bailey et al. (1963).
2. The resulting squared errors from the regression in the first step are regressed on a constant term and on the time between sales (holding period) as follows

$$\mathbf{U} = \boldsymbol{\beta}_0 + \boldsymbol{\beta}_1 \mathbf{HP} + \mathbf{H}. \quad (4.4)$$

⁴⁰See <https://www.ft.com/content/f36d7eeb-03d6-4dab-909a-d5d4524a4edb>.

⁴¹Note that we stated this effect to be negligible in our setting.

$\mathbf{HP}$ is the vector including the holding periods (in weeks) and $\mathbf{H}$ is the vector with the residuals, which are assumed to have an expected value of zero ($E[\mathbf{H}] = 0$). The reciprocal of the square root of the fitted values from the regression are the weights (w_i). We incorporate them to the weighting vector mentioned above to obtain the final weighting vector $\mathbf{\Delta} = [w_1 * \omega_1 \quad w_2 * \omega_2 \quad \dots \quad w_N * \omega_N]$. The weighting should reduce the effect of sale pairs with large gap times on the index estimates and treat virtual land plots with more than two sales appropriately.

3. The weighted least square regression is run by repeating the first step but incorporating the weights from the second step to obtain the final weighted regression

$$\mathbf{Y} = \mathbf{\Delta X} \boldsymbol{\beta} + \mathbf{U}. \quad (4.5)$$

The index values are then again retrieved by exponentiating the estimated coefficients $\hat{\beta}_t$.

A.4 Wavelet analysis – Methodology

Wavelet analysis starts with the wavelet transform, which refers to the process of decomposing a time series into its individual elementary functions, also referred to as daughter wavelets or simply wavelets. These wavelets capture specific features of the time series and are translated and scaled versions of a mother wavelet $\psi(t)$. They are defined as

$$\psi_{(\tau,s)}(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t-\tau}{s}\right), \quad (4.6)$$

where τ is the time position (translation parameter) and s is the scale (dilation parameter), which represents the frequency. The frequency indicates how often the wavelet function oscillates within a given time frame. Higher frequencies accord with more rapid oscillations and thus capture smaller details in the time series, while lower frequencies accord with slower oscillations and capture broader trends. There exists a range of different mother wavelets for various time series characteristics. The most popular mother wavelet used in the field of economic and finance is the Morlet wavelet (see, e.g., Rua and Nunes, 2009) due to its optimal joint time-frequency concentration. In other words, the Morlet wavelet can effectively address the fundamental trade-off problem of simultaneous time and frequency resolution in the context of time-frequency analysis.

A mother wavelet must satisfy several properties. First, its value must integrate to zero: $\int_{-\infty}^{+\infty} \psi(t) dt = 0$. Second, its absolute value must integrate to one: $\int_{-\infty}^{+\infty} |\psi(t)| dt = 1$, which implies that $\psi(t)$ is restricted to a finite interval of time. Lastly, it satisfies the admissibility condition given by

$$C_{\text{adm},\psi} = \int_0^{+\infty} \frac{|\hat{\psi}(\omega)|^2}{\omega} d\omega < \infty, \quad (4.7)$$

where $\hat{\psi}(\omega)$ is the Fourier transform of $\psi(t)$, defined as $\hat{\psi}(\omega) = \int_{-\infty}^{+\infty} \psi(t) e^{-i\omega t} dt$.⁴² By satisfying the admissibility condition, we can recover a time series $x(t)$ from its continuous wavelet transform $W_x(\tau, s)$, which can be defined as

$$W_x(\tau, s) = \int_{-\infty}^{+\infty} x(t) \psi_{(\tau,s)}^*(t) dt = \frac{1}{\sqrt{s}} \int_{-\infty}^{+\infty} x(t) \psi\left(\frac{t-\tau}{s}\right) dt, \quad (4.8)$$

where $*$ denotes the complex conjugate.

We follow the common literature (see, e.g., Rua and Nunes, 2009) and use the Morlet wavelet as the mother wavelet for our analysis. The Morlet wavelet is a complex sine wave within a Gaussian envelope and is defined as

$$\psi(t) = \pi^{-\frac{1}{4}} e^{i\omega_0 t} e^{-\frac{t^2}{2}}. \quad (4.9)$$

⁴²The Fourier transform decomposes a signal or time series in the time domain into its frequency component and their corresponding amplitudes, but does not provide information about when specific frequencies occur like the wavelet transform does. For a more detailed explanation on the Fourier transform and wavelet transform, see, e.g., Torrence and Compo (1998).

where ω_0 is the wavenumber of the wavelet, which is often set to 6 as it provides a good balance between time and frequency (see, e.g., Rua and Nunes, 2009; Torrence and Compo, 1998). The corresponding Fourier transform is then given by $\hat{\psi}(\omega) = \pi^{\frac{1}{4}} \sqrt{2} e^{-\frac{1}{2}(\omega - \omega_0)^2}$.

Now, we calculate the cross-wavelet spectrum, which measures the relationship between two time series $x(t)$ and $y(t)$ in the time-frequency space with their wavelet transforms $W_x(\tau, s)$ and $W_y(\tau, s)$. The cross-wavelet spectrum can be denoted as

$$W_{xy}(\tau, s) = W_x(\tau, s) * W_y^*(\tau, s), \quad (4.10)$$

which is the product of the wavelet transform of $x(t)$ at a given time scale (s) and position (τ) by the complex conjugate, represented by $(*)$ of the wavelet transform of $y(t)$ at the same time scale and position.

With the cross-wavelet spectrum we are able to calculate the wavelet squared coherency, which measures the similarity or coherence to which two time series move together around each moment in time and for each frequency (time-frequency domain). Its intuition can be interpreted like the squared coefficient of correlation $R^2(\tau, s)$, since it also ranges between 0 (no co-movement or correlation) and 1 (perfect co-movement or correlation). But, the wavelet approach allows for the detection of both time-varying and frequency-varying features, providing a more comprehensive understanding of the relationship between the two time series. Following Torrence and Compo (1998), we calculate the cross-wavelet spectrum by taking the absolute value squared of the smoothed cross-wavelet spectrum $W_{xy}(\tau, s)$ and normalizing it by the product of the absolute value squared of the smoothed wavelet transforms $W_x(\tau, s)$ and $W_y(\tau, s)$ as follows⁴³


$$R^2(\tau, s) = \frac{|S(s^{-1}W_{xy}(\tau, s))|^2}{S(s^{-1}|W_x(\tau, s)|^2)S(s^{-1}|W_y(\tau, s)|^2)}. \quad (4.11)$$

By analyzing the graph of the wavelet squared coherency (scalogram), one can identify the specific regions in the time-frequency domain where the two time series exhibit co-movement.

⁴³The smoothing operation $S(\cdot)$ is necessary to reduce noise and enhance the underlying trends and coherent patterns of the time series. For more details, see, e.g., Grinsted et al. (2004).


A.5 Wavelet analysis – Further results

Figure A.4: Virtual land & FTSE Nareit All Equity REITs Index




Source(s): Authors' own work based on data from Opensea and Nareit.

Figure A.5: Virtual land & FTSE Nareit All Equity REITs Index




Source(s): Authors' own work based on data from Opensea and Nareit.

Figure A.6: Virtual land & S&P Global REIT Index



Source(s): Authors' own work based on data from Opensea and S&P Global.

Figure A.7: Virtual land & S&P Global REIT Index




Source(s): Authors' own work based on data from Opensea and S&P Global.

A.6 Kendall's Tau

Kendall's τ correlation measure assesses the ordinal association between the virtual land market and physical real estate market returns by evaluating the ranks and the concordance of pairs, making it robust to outliers and suitable for non-linear relationships. The results are very similar to the Pearson correlation measure.

Figure A.8: Virtual land & REITs (Kendall's τ)



Source(s): Authors' own work based on data from Opensea, Nareit, and S&P Global.

A.7 Description of variables

Table A.6: Description of variables

Variable name	Description	Data source / paper
Macroeconomic data		
Consumer Sentiment Index	Survey on Consumer Sentiment across the USA	Surveys of Consumers
Economic Conditions Index	Survey on Economic Conditions across the USA	Surveys of Consumers
Global Economic Policy Uncertainty Index	Uncertainty based on news data	Baker et al. (2016)
US Economic Policy Uncertainty Index	Uncertainty based on news data	Baker et al. (2016)
Δ Fed Funds Effective Rate	Monthly change of Fed Funds Rate	FRED
US M2 money supply (Bn USD)	M2 money supply in the USA	FRED
VIX Index	CBOE Volatility Index	FRED
Aruoba-Diebold-Scotti Index	Business conditions index to track real business conditions	Aruoba et al. (2009)
Termspread	10Y-3M Treasury yield spread	FRED
Crypto Data		
MANA (USD)	Price of Decentraland's cryptocurrency in USD	CoinMarketCap
ETH (USD)	Price of Ethereum's cryptocurrency in USD	CoinMarketCap
ETH Gas Fees (USD)	Price of a transaction on the Ethereum blockchain	Etherscan
Volatility of MANA	30-day volatility of Decentraland's cryptocurrency in USD	CoinMarketCap
Volatility of ETH	30-day volatility of Ethereum's cryptocurrency in USD	CoinMarketCap
Crypto Policy Uncertainty Index	Uncertainty around crypto policy based on news	Lucey et al. (2022)
Crypto Price Uncertainty Index	Uncertainty around crypto prices based on news	Lucey et al. (2022)
Sentiment Data		
Google Search "Metaverse Real Estate"	Google Search volume of the requested words	Google Trends
Google Search "Buy Real Estate"	Google Search volume of the requested words	Google Trends
Google Search "Virtual Land"	Google Search volume of the requested words	Google Trends
Google Search "Metaverse"	Google Search volume of the requested words	Google Trends

Source(s): Authors' own work.

A.8 Portfolio analysis

Table A.7: Performance metrics by year and asset (weekly basis)

Year	Asset	μ	SD	Sharpe	Sortino
2019	Virtual Land	<i>0.0064</i>	0.1770	0.0342	0.0406
	REITs	0.0034	0.0132	0.2291	0.2074
	Commercial REITs	0.0023	0.0179	0.1051	0.0944
	Stocks	0.0043	0.0139	<i>0.2775</i>	<i>0.2916</i>
	Bonds	0.0013	<i>0.0053</i>	0.1739	0.1596
2020	Virtual Land	<i>0.0319</i>	0.2449	0.1299	<i>0.1665</i>
	REITs	-0.0001	0.0667	-0.0021	-0.0019
	Commercial REITs	-0.0007	0.0937	-0.0080	-0.0081
	Stocks	0.0034	0.0432	0.0777	0.0676
	Bonds	0.0017	<i>0.0109</i>	<i>0.1524</i>	0.1328
2021	Virtual Land	<i>0.0790</i>	0.2311	<i>0.3421</i>	<i>0.5541</i>
	REITs	0.0049	0.0158	0.3092	0.3235
	Commercial REITs	0.0061	0.0237	0.2556	0.2518
	Stocks	0.0037	0.0163	0.2239	0.2270
	Bonds	-0.0011	<i>0.0044</i>	-0.2434	-0.2056
2022	Virtual Land	-0.0383	0.1383	-0.2794	-0.2579
	REITs	-0.0055	0.0300	-0.1946	-0.2053
	Commercial REITs	-0.0049	0.0328	-0.1596	-0.1568
	Stocks	-0.0037	0.0297	<i>-0.1382</i>	<i>-0.1554</i>
	Bonds	<i>-0.0035</i>	<i>0.0124</i>	-0.3175	-0.3138
2023	Virtual Land	<i>0.0236</i>	0.2093	0.1082	0.1282
	REITs	0.0015	0.0274	0.0199	0.0216
	Commercial REITs	0.0007	0.0357	-0.0062	-0.0072
	Stocks	0.0039	0.0189	<i>0.1579</i>	<i>0.1707</i>
	Bonds	0.0012	<i>0.0105</i>	0.0258	0.0325
Total	Virtual Land	<i>0.0205</i>	0.2057	<i>0.0980</i>	<i>0.1240</i>
	REITs	0.0009	0.0361	0.0139	0.0129
	Commercial REITs	0.0007	0.0488	0.0072	0.0073
	Stocks	0.0023	0.0267	0.0731	0.0697
	Bonds	-0.0001	<i>0.0094</i>	-0.0458	-0.0451

Note: This table shows the expected returns (μ), the standard deviations (SD), the Sharpe ratios and the Sortino ratios by year and by asset, as well as for the entire period (Total). REITs are represented by the S&P Global REIT Index (USD). Commercial REITs are represented by our constructed FTSE Nareit All Equity REITs Index (only Retail/Office/Gaming) (USD). Stocks are represented by the MSCI World Index (USD). Bonds are represented by the FTSE World Broad Investment-Grade Bond Index (USD). The expected returns and the standard deviations are calculated using weekly returns. The ratios are calculated using weekly excess returns over the risk-free rate, which is the 3-Month Treasury Bill Secondary Market Rate (<https://fred.stlouisfed.org/series/DTB3#0>). Italic numbers highlight the best performance for the respective metric. Source(s): Authors' own work based on data from Opensea, Nareit, S&P Global, FTSE, MSCI, and FRED.

Table A.8: Performance metrics by year and asset (annualized)

Year	Asset	μ	SD	Sharpe	Sortino
2019	Virtual Land	<i>0.33</i>	1.28	0.25	0.29
	REITs	0.18	0.10	1.65	1.50
	Commercial REITs	0.12	0.13	0.76	0.68
	Stocks	0.22	0.10	<i>2.00</i>	<i>2.10</i>
	Bonds	0.07	<i>0.04</i>	1.25	1.15
2020	Virtual Land	<i>1.66</i>	1.77	0.94	<i>1.20</i>
	REITs	-0.00	0.48	-0.01	-0.01
	Commercial REITs	-0.04	0.68	-0.06	-0.06
	Stocks	0.18	0.31	0.56	0.49
	Bonds	0.09	<i>0.08</i>	<i>1.10</i>	0.96
2021	Virtual Land	<i>4.11</i>	1.67	<i>2.47</i>	<i>4.00</i>
	REITs	0.26	0.11	2.23	2.33
	Commercial REITs	0.32	0.17	1.84	1.82
	Stocks	0.19	0.12	1.61	1.64
	Bonds	-0.06	<i>0.03</i>	-1.76	-1.48
2022	Virtual Land	-1.99	1.00	-2.01	-1.86
	REITs	-0.28	0.22	-1.40	-1.48
	Commercial REITs	-0.25	0.24	-1.15	-1.13
	Stocks	-0.19	0.21	<i>-1.00</i>	<i>-1.12</i>
	Bonds	<i>-0.18</i>	<i>0.09</i>	-2.29	-2.26
2023	Virtual Land	<i>1.23</i>	1.51	0.78	0.92
	REITs	0.08	0.20	0.14	0.16
	Commercial REITs	0.04	0.26	-0.04	-0.05
	Stocks	0.20	0.14	<i>1.14</i>	<i>1.23</i>
	Bonds	0.06	<i>0.08</i>	0.19	0.23
Total	Virtual Land	<i>1.07</i>	1.48	<i>0.71</i>	<i>0.89</i>
	REITs	0.04	0.26	0.10	0.09
	Commercial REITs	0.04	0.35	0.05	0.05
	Stocks	0.12	0.19	0.53	0.50
	Bonds	-0.00	<i>0.07</i>	-0.33	-0.32

Note: This table shows the annualized expected returns (μ), the standard deviations (SD), the Sharpe ratios and the Sortino ratios by year and by asset, as well as for the entire period (Total). REITs are represented by the S&P Global REIT Index (USD). Commercial REITs are represented by our constructed FTSE Nareit All Equity REITs Index (only Retail/Office/Gaming) (USD). Stocks are represented by the MSCI World Index (USD). Bonds are represented by the FTSE World Broad Investment-Grade Bond Index (USD). The expected returns and the standard deviations are calculated using weekly returns. The ratios are calculated using weekly excess returns over the risk-free rate, which is the 3-Month Treasury Bill Secondary Market Rate (<https://fred.stlouisfed.org/series/DTB3#0>). Italic numbers highlight the best performance for the respective metric. Source(s): Authors' own work based on data from Opensea, Nareit, S&P Global, FTSE, MSCI, and FRED.

Table A.9: Correlation matrix of excess returns

	Virtual land	REITs	Commercial REITs	Stocks	Bonds
Virtual land	1.00	0.14	0.08	0.22	0.13
REITs		1.00	0.93	0.85	0.56
Commercial REITs			1.00	0.74	0.40
Stocks				1.00	0.49
Bonds					1.00

Note: This table shows the correlation matrix of excess returns for the selected assets over the sample period from January 2019 to December 2023. REITs are represented by the S&P Global REIT Index (USD). Commercial REITs are represented by our constructed FTSE Nareit All Equity REITs Index (only Retail/Office/Gaming) (USD). Stocks are represented by the MSCI World Index (USD). Bonds are represented by the FTSE World Broad Investment-Grade Bond Index (USD). The excess returns are calculated by subtracting the weekly risk-free rate from the asset returns. Source(s): Authors' own work based on data from Opensea, Nareit, S&P Global, FTSE, MSCI, and FRED.

Table A.10: Portfolio weights

Optimization	Portfolio	REITs	Commercial REITs	Stocks	Bonds	Weight Virtual land
1/n	Benchmark	0.25	0.25	0.25	0.25	
	(+Virtual land)	0.20	0.20	0.20	0.20	0.20
Max Sharpe	Benchmark	0.00	0.00	1.00	0.00	
	(+Virtual land)	0.00	0.00	0.83	0.00	0.17
Max Sortino	Benchmark	0.00	0.00	1.00	0.00	
	(+Virtual land)	0.00	0.00	0.60	0.00	0.40
Min variance	Benchmark	0.15	0.00	0.30	0.55	
	(+Virtual land)	0.15	0.00	0.30	0.55	0.00
REITs	Benchmark	1.00	0.00	0.00	0.00	
	(+ 10% Virtual land)	0.90	0.00	0.00	0.00	0.10
	(+ 20% Virtual land)	0.80	0.00	0.00	0.00	0.20
Commercial REITs	Benchmark	1.00	0.00	0.00	0.00	
	(+ 10% Virtual land)	0.90	0.00	0.00	0.00	0.10
	(+ 20% Virtual land)	0.80	0.00	0.00	0.00	0.20
Multi asset	Benchmark	0.10	0.00	0.70	0.20	
	(+ 10% Virtual land)	0.10	0.00	0.60	0.20	0.10

Note: This table shows the portfolio weights for each asset in each portfolio strategy corresponding to the portfolio performance comparison in Table 4.5. Source(s): Authors' own work.

References

- Aharon D. Y. and Demir E. (2022). NFTs and asset class spillovers: Lessons from the period around the COVID-19 pandemic. *Finance Research Letters* 47, p. 102515.
- Ante L. (2022). The non-fungible token (NFT) market and its relationship with Bitcoin and Ethereum. *FinTech* 1.3, pp. 216–224.
- Ante L., Wazinski F.-P., and Saggu A. (2023). Digital real estate in the metaverse: An empirical analysis of retail investor motivations. *Finance Research Letters* 58, p. 104299.
- Aruoba S. B., Diebold F. X., and Scotti C. (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics* 27.4, pp. 417–427.
- Bailey M. J., Muth R. F., and Nourse H. O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association* 58.304, pp. 933–942.
- Baker S. R., Bloom N., and Davis S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131.4, pp. 1593–1636.
- Baur D. G. and McDermott T. K. (2010). Is gold a safe haven? International evidence. *Journal of Banking and Finance* 34.8, pp. 1886–1898.
- Beaupain R. and Heck S. (2016). A repeat-sales index for pricing US corporate bonds. *Finance* 37.2, pp. 75–117.
- Borri N., Liu Y., and Tsyvinski A. (2023). The economics of non-fungible tokens. SSRN. Available at URL: <https://ssrn.com/abstract=4052045>.
- Boudry W. I., Coulson N. E., Kallberg J. G., and Liu C. H. (2012). On the hybrid nature of REITs. *The Journal of Real Estate Finance and Economics* 44.1, pp. 230–249.
- Buterin V. (2014). Ethereum: A next-generation smart contract and decentralized application platform. White Paper. URL: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- Case K. E. and Shiller R. J. (1987). Prices of single-family homes since 1970: New indexes for four cities. *New England Economic Review* September/October, pp. 45–56.
- Christoffersen P., Errunza V., Jacobs K., and Langlois H. (2012). Is the potential for international diversification disappearing? A dynamic copula approach. *The Review of Financial Studies* 25.12, pp. 3711–3751.
- Clayton J. and MacKinnon G. (2003). The relative importance of stock, bond and real estate factors in explaining REIT returns. *The Journal of Real Estate Finance and Economics* 27.1, pp. 39–60.
- CoinMarketCap (2024). Historical data of ETH, WETH, and MANA, 1 January 2019 to 31 December 2023. Accessed 13 March 2024. Data retrieved from <https://coinmarketcap.com/currencies/ethereum/historical-data/>, <https://coinmarketcap.com/currencies/weth/historical-data/>, and <https://coinmarketcap.com/currencies/decentraland/>.
- DeMiguel V., Garlappi L., and Uppal R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies* 22.5, pp. 1915–1953.
- DiPasquale D. and Wheaton W. C. (1992). The markets for real estate assets and space: A conceptual framework. *Real Estate Economics* 20.2, pp. 181–198.
- Dowling M. (2022). Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters* 44, p. 102097.

- Fenn J. and Raskino M. (2007). Understanding Gartner's hype cycles, 2007. Gartner ID G00144727. Accessed 19 March 2024. URL: <https://www.gartner.com/en/documents/509085>.
- Ghosh I., Alfaro-Cortés E., Gámez M., and García N. (2023). Do travel uncertainty and invasion rhetoric spur metaverse financial assets? Gauging the role of media influence. *Finance Research Letters* 51, p. 103434.
- Giang Barrera K. and Shah D. (2023). Marketing in the metaverse: Conceptual understanding, framework, and research agenda. *Journal of Business Research* 155, p. 113420.
- Goanta C. (2020). Selling LAND in Decentraland: The regime of non-fungible tokens on the Ethereum blockchain under the digital content directive. *Disruptive technology, legal innovation, and the future of real estate*. Ed. by A. Lehavi and R. Levine-Schnur. Springer, pp. 139–154.
- Goldberg M. and Schär F. (2023). Metaverse governance: An empirical analysis of voting within decentralized autonomous organizations. *Journal of Business Research* 160, p. 113754.
- Goldberg M., Schär F., and Kugler P. (2024). Land valuation in the metaverse: Location matters. *Journal of Economic Geography* 24.1, pp. 67–92.
- Goldberg M., Schär F., and Thürkauf D. (2023). Retailing and customer engagement in the metaverse: An empirical analysis. SSRN. Available at URL: <https://ssrn.com/abstract=4547100>.
- Grinsted A., Moore J. C., and Jevrejeva S. (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics* 11.5/6, pp. 561–566.
- Haurin D. and Hendershott P. (1991). House price indexes: Issues and results. *Real Estate Economics* 19.3, pp. 259–269.
- Heikkila E., Gordon P., Kim J. I., Peiser R. B., Richardson H. W., and Dale-Johnson D. (1989). What happened to the CBD-distance gradient? Land values in a policentric city. *Environment and Planning A: Economy and Space* 21.2, pp. 221–232.
- Hoesli M. and Oikarinen E. (2012). Are REITs real estate? Evidence from international sector level data. *Journal of International Money and Finance* 31.7, pp. 1823–1850.
- Hoesli M. and Oikarinen E. (2016). Are public and private asset returns and risks the same? Evidence from real estate data. *Journal of Real Estate Portfolio Management* 22.2, pp. 179–198.
- Hoesli M. and Oikarinen E. (2021). Does listed real estate behave like direct real estate? Updated and broader evidence. *Applied Economics* 53.26, pp. 3023–3042.
- Hollensen S., Kotler P., and Opresnik M. O. (2022). Metaverse – The new marketing universe. *Journal of Business Strategy* 44.3, pp. 119–125.
- Hudson-Smith A. (2022). Incoming metaverses: Digital mirrors for urban planning. *Urban Planning* 7.2.
- Humphreys B. R., Nowak A., and Zhou Y. (2019). Superstition and real estate prices: Transaction-level evidence from the US housing market. *Applied Economics* 51.26, pp. 2818–2841.
- Jain A., Behzad S., Noone C., Garlichs I., Brophy T., and McMullen P. (2022). Metaverse – The next e-commerce revolution. Deutsche Bank Research. Accessed 19 March 2024. URL: <https://corporates.db.com/publications/White-papers-guides/metaverse-the-next-e-commerce-revolution>.
- Knowles I. and Castronova E. (2016). The economics of virtual worlds. *Handbook on the Economics of the Internet*. Edward Elgar Publishing, pp. 495–509.

- Ko H. and Lee J. (2023). Non-fungible tokens: A hedge or a safe haven? *Applied Economics Letters*, pp. 1–8.
- Ko H., Son B., Lee Y., Jang H., and Lee J. (2022). The economic value of NFT: Evidence from a portfolio analysis using mean–variance framework. *Finance Research Letters* 47, p. 102784.
- Kong D.-R. and Lin T. C. (2021). Alternative investments in the fintech era: The risk and return of non-fungible token (NFT). SSRN. Available at URL: <https://ssrn.com/abstract=3914085>.
- Kreppmeier J., Laschinger R., Steininger B. I., and Dorfleitner G. (2023). Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *Journal of Banking and Finance* 154, p. 106940.
- Kumar A. S. and Padakandla S. R. (2023). Do NFTs act as a good hedge and safe haven against cryptocurrency fluctuations? *Finance Research Letters* 56, p. 104131.
- Lindenthal T. (2014). Valuable words: The price dynamics of internet domain names. *Journal of the Association for Information Science and Technology* 65.5, pp. 869–881.
- Lindenthal T. (2018). Monocentric cyberspace: The primary market for internet domain names. *The Journal of Real Estate Finance and Economics* 57.1, pp. 152–166.
- Ling D. C. and Naranjo A. (1999). The integration of commercial real estate markets and stock markets. *Real Estate Economics* 27.3, pp. 483–515.
- Lucey B. M., Vigne S. A., Yarovaya L., and Wang Y. (2022). The cryptocurrency uncertainty index. *Finance Research Letters* 45, p. 102147.
- Maouchi Y., Charfeddine L., and El Montasser G. (2022). Understanding digital bubbles amidst the COVID-19 pandemic: Evidence from DeFi and NFTs. *Finance Research Letters* 47, p. 102584.
- Markowitz H. (1952). Portfolio selection. *The Journal of Finance* 7.1, pp. 77–91.
- Mei J. and Moses M. (2002). Art as an investment and the underperformance of masterpieces. *The American Economic Review* 92.5, pp. 1656–1668. URL: <https://www.jstor.org/stable/3083271>.
- Momtaz P. P. (2022). Some very simple economics of Web3 and the metaverse. *FinTech* 1.3, pp. 225–234.
- Mou J., Liu W., Guan C., Westland J. C., and Kim J. (2024). Predicting the cryptocurrency market using social media metrics and search trends during COVID-19. *Electronic Commerce Research* 24.2, pp. 1307–1333.
- Nagl C., Nagl M., Rösch D., Schäfers W., and Freybote J. (2024). Time varying dependences between real estate crypto, real estate and crypto returns. *Journal of Real Estate Research* 46.4, pp. 538–566.
- Nakavachara V. and Saengchote K. (2022). Does unit of account affect willingness to pay? Evidence from metaverse LAND transactions. *Finance Research Letters* 49, p. 103089.
- Newey W. K. and West K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review* 28.3, pp. 777–787.
- OneLand (2024). Historical snapshot of virtual land market capitalization – 19 March 2024. Accessed 19 March 2024. URL: <https://oneland.world/#/main>.
- Ooi K.-B., Tan G. W.-H., Aw E. C.-X., Cham T.-H., Dwivedi Y. K., Dwivedi R., Hughes L., Kar A. K., Loh X.-M., Mogaji E., Phau I., and Sharma A. (2023). Banking in the metaverse: A new frontier for financial institutions. *International Journal of Bank Marketing* 41.7, pp. 1829–1846.

- Ordano E., Meilich A., Jardi Y., and Araoz M. (2017). Decentraland: A blockchain-based virtual world. White Paper. URL: <https://decentraland.org/whitepaper.pdf>.
- Pagliari J. L., Scherer K. A., and Monopoli R. T. (2005). Public versus private real estate equities: A more refined, long-term comparison. *Real Estate Economics* 33.1, pp. 147–187.
- Pesando J. E. (1993). Art as an investment: The market for modern prints. *The American Economic Review* 83.5, pp. 1075–1089. URL: <https://www.jstor.org/stable/2117549>.
- Proelss J., Sévigny S., and Schweizer D. (2023). GameFi: The perfect symbiosis of blockchain, tokens, DeFi, and NFTs? *International Review of Financial Analysis* 90, p. 102916.
- Rua A. and Nunes L. C. (2009). International comovement of stock market returns: A wavelet analysis. *Journal of Empirical Finance* 16.4, pp. 632–639.
- Saengchote K., Nakavachara V., and Xu Y. (2023). Capitalising the network externalities of new land supply in the metaverse. SSRN. Available at URL: <https://ssrn.com/abstract=4404556>.
- Savills (2023). Total value of global real estate. Accessed 19 March 2024. URL: <https://www.savills.com/impacts/market-trends/the-total-value-of-global-real-estate-property-remains-the-worlds-biggest-store-of-wealth.html>.
- Schär F. and Goldberg M. (2023). Scarcity in the metaverse: Space, location and the attention economy. SSRN. Available at URL: <https://ssrn.com/abstract=4606813>.
- Schumacher P. (2022). The metaverse as opportunity for architecture and society: Design drivers, core competencies. *Architectural Intelligence* 1.11.
- Serneels S. (2023). Detecting wash trading for nonfungible tokens. *Finance Research Letters* 52, p. 103374.
- Sharpe W. F. (1994). The Sharpe ratio. *The Journal of Portfolio Management* 21.1, pp. 49–58.
- Shaver G., Del Guidice A., Chen J., Qiu Tiantian G., and Huang Peng H. (2023). The role of REITs in real estate allocations. Accessed 19 March 2024. URL: https://www.gic.com.sg/wp-content/uploads/2023/03/GIC-ThinkSpace-The-Role-of-REITs-in-Real-Estate-Allocations_v1.pdf.
- Shiller R. J. (1991). Arithmetic repeat sales price estimators. *Journal of Housing Economics* 1.1, pp. 110–126.
- Simons R. A. and Simons S. T. (2022). Blockchain, cryptocurrency, and real estate: The current situation and prospects for the next 5 years. *Journal of Real Estate Practice and Education* 24.1, pp. 72–85.
- Sortino F. A. and Meer R. van der (1991). Downside risk: Capturing what’s at stake in investment situations. *Journal of Portfolio Management* 17.4, pp. 27–31.
- Swinkels L. (2023). Empirical evidence on the ownership and liquidity of real estate tokens. *Financial Innovation* 9, pp. 1–29.
- Torrence C. and Compo G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society* 79.1, pp. 61–78.
- Umar Z., Gubareva M., Teplova T., and Tran D. K. (2022). Covid-19 impact on NFTs and major asset classes interrelations: Insights from the wavelet coherence analysis. *Finance Research Letters* 47, p. 102725.
- Urom C., Ndubuisi G., and Guesmi K. (2022). Dynamic dependence and predictability between volume and return of non-fungible tokens (NFTs): The roles of market factors and geopolitical risks. *Finance Research Letters* 50, p. 103188.

- Vidal-Tomás D. (2022). The new crypto niche: NFTs, play-to-earn, and metaverse tokens. *Finance Research Letters* 47, p. 102742.
- Wang X.-l., Ye Q., and Lu Q. (2010). The influence of real world on virtual world: An investigation of virtual land price in Second Life. *2010 International Conference on Management Science & Engineering: 17th Annual Conference Proceedings*, pp. 625–630.
- Westcott K., Arbanas J., Arkenberg C., and Auxier B. (2023). 2023 digital media trends: Immersed and connected. Deloitte Insights. Accessed 2 March 2026. URL: <https://www2.deloitte.com/us/en/insights/industry/technology/digital-media-trends-consumption-habits-survey/2023.html#read-the-report>.
- Wood G. (2014). Ethereum: A secure decentralised generalised transaction ledger. Yellow Paper. URL: <https://ethereum.github.io/yellowpaper/paper.pdf>.
- Yencha C. (2023). Spatial heterogeneity and non-fungible token sales: Evidence from Decentraland LAND sales. *Finance Research Letters* 58, p. 103628.
- Yi R. Y. (2019). Decentraland lords. Accessed 19 March 2024. URL: <https://tokeneconomy.co/decentraland-lords-9953d0de1e5c>.
- Yousaf I. and Yarovaya L. (2022). Static and dynamic connectedness between NFTs, DeFi and other assets: Portfolio implication. *Global Financial Journal* 53, p. 100719.

Conflicts of Interest

One author bought a small amount of Decentraland's virtual currency token (MANA) to better understand the platform ecosystem and the functioning of its marketplace. The current value of these holdings is less than \$100. The authors have no commercial relationship with Decentraland, whose data this study primarily relies on, or with any other virtual land platform.

Acknowledgments

We would like to express our gratitude to Editor Olayiwola Oladiran and the anonymous reviewers at the *Journal of European Real Estate Research* for their insightful comments and suggestions, which have greatly enhanced the quality of this paper. Furthermore, we would like to thank Fabian Schär (Center for Innovative Finance, University of Basel) for proofreading the manuscript and for his valuable feedback. We also appreciate the suggestions from the participants at the 2023 European Real Estate Society Annual Conference in London, UK, and the 2024 American Real Estate Society Annual Conference in Orlando, USA.

Chapter 5

Conclusion

The final chapter of this dissertation provides an executive summary of each paper presented in the preceding chapters, encapsulating the identified problem and research objective, the curated data, and the methodological approach and econometric techniques employed. The results and implications form the heart of the summaries, providing practitioners, the scientific community, policymakers, and regulators with the key takeaways from the research. It concludes with final remarks and an outlook on a possible path forward for real estate markets in the blockchain era.

5.1 Executive Summary

Paper 1: Liquidity Mechanisms in Decentralized Finance: Design and Fragmentation in Illiquid Real-World Asset Markets

Problems and objective

The tokenization of illiquid real-world assets (RWAs), such as real estate, is a promising new innovation that alters market structure not only in terms of how assets are represented and fractionalized, but also in how trading is organized, and market liquidity is created, allocated, and preserved (Whitaker and Kräussl, 2020). Conceptually, lumpy single assets such as real estate or a piece of art become divisible and more tradable, similar to stocks or fund shares. While the primary token market for these formerly illiquid RWAs has already been studied more extensively (see, e.g., Swinkels, 2023; Kreppmeier et al., 2023; Bergkamp et al., 2025), the secondary market and how liquidity is formed and trading is organized remain largely unexplored. Decentralized finance (DeFi), a broader term for an open blockchain-based financial market system, promises a disintermediated form of liquidity transformation or, put differently, the creation of endogenous liquidity through the design of market mechanisms in a new blockchain-based guise (Schär, 2021). While DeFi is still largely dominated by cryptocurrencies, tokenized RWAs are increasingly integrated into it, testing its techno-economic boundaries (Cong et al., 2025b). This paper explores three prominent blockchain-based liquidity mechanisms for tokenized illiquid RWAs, each characterized by distinct pricing

rules, liquidity provision, and execution certainty. While tokenization facilitates secondary markets in general, the central question this paper aims to answer is how the coexistence of heterogeneous liquidity mechanisms shapes trading and liquidity formation for tokenized illiquid RWAs.

Data and methodology

This paper studies one of the earliest and largest tokenized real estate ecosystems, RealT, which tokenizes rental residential real estate in the United States and Latin America. From 2019 to October 2025, which defines our observation window, the platform has tokenized approximately \$163 million worth of these assets. Its secondary market provides a rare empirical setting in which three distinct liquidity mechanisms coexist on public blockchain infrastructure, allowing access and direct measurement of mechanism-specific trading and liquidity outcomes. The study follows a three-part methodological approach. First, a self-contained theoretical framework of trading execution costs and liquidity provision is presented. Second, the paper reconstructs and processes data from the Ethereum and Gnosis blockchain, where all transactions in the sample are recorded. The resulting dataset comprises 777,038 trades, \$41.2 million in traded volume, 9,202 unique wallet addresses, and 752 tokenized real estate assets across various transaction-level and liquidity-state panels. Third, this on-chain microdata is examined using ordinary least squares (OLS) panel regressions complemented by difference-in-differences analyses around two exogenous shocks, one DeFi-wide and one platform-specific, to ensure causal identification.

Results and their contribution to science and practice

With its unique composite dataset on three coexisting liquidity mechanisms, the study provides the first targeted and comprehensive empirical analyses of whether tokenization can endogenously create and sustain liquidity for formerly hard-to-trade assets. The main finding is that, while tokenization can catalyze secondary-market formation for illiquid RWAs, it does not per se ensure liquidity. While the lot-size transformation of lumpy assets enabled through tokenization can lower investment barriers in the primary issuance (Swinkels, 2023; Kreppmeier et al., 2023), the small capitalization of the underlying asset still limits trading demand and external liquidity provision in the secondary market, which makes incentive compatibility, mechanism design, and careful orchestration of coexisting execution options essential. The study formally and empirically analyzes these frictions and demonstrates that automated market makers (AMMs), although very successful in the trading of liquid cryptocurrencies (Barbon and Rinaldo, 2026), struggle to scale for small, slow-moving, and heterogeneous RWAs when trading activity is limited. In such settings, the asset's liquidity pool remains limited by the size of the asset itself and is generally too thin, making pool-based price impacts unbearable and trading-fee-based liquidity provision unprofitable. For the AMM mechanism to work more effectively, the assets to be tokenized may need to be larger or pooled together to attract more trading demand and liquidity, or the incentive mechanisms for external liquidity provision may need to be rethought. Peer-to-peer (P2P) trading dominates volume in the studied setting, but ultimately resembles a digitized search market without incentives to provide continuous liquidity in its current form. The Buyback option acts as an important liquidity backstop, whose suspension has negative effects on overall market liquidity, but does not endogenously create liquidity because it is provided by the issuer. The results carry novel and timely implications for science and practice. The derived framework can be extended to other and future combinations of coexisting market mechanisms in tokenized illiquid RWA

markets, and the rich micro-level data and curation procedure can serve as a foundational dataset for subsequent empirical studies in this field. For practice, tokenization platforms and exchange protocol designers need to ensure sufficient trading demand in order to attract external liquidity provision or rethink their market design structures. Investment managers contemplating tokenization as an alternative investment vehicle can draw on this empirical evidence when considering which assets to tokenize and which market mechanisms to pursue. Lastly, the results can help regulators and policymakers make informed decisions on potential redemption rules and other market interventions to protect market participants without hindering innovation.

Paper 2: Tokenization, Transparency, and Time on Market: Empirical Evidence from a Blockchain-Based Real Estate Marketplace

Problems and objective

The limited transparency over the listing and selling histories in real estate markets complicates not only the exchange of real estate itself, but also restricts the ability to study the relationship between pricing behaviour and Time on Market (ToM), a common measure of liquidity in real estate markets. Blockchain technology moves traditionally opaque transaction and listing histories into a digitally recorded environment in which listing, repricing, and execution behavior can be observed far more closely than in conventional property markets. More precisely, conventional real estate datasets usually contain only the initial listing price and the final selling price, while intermediate price revisions are rarely observed, with only a few exceptions (see, e.g., Knight, 2002; Merlo and Ortalo-Magné, 2004). Against this background, the objective of this study is to examine how blockchain-enabled traceability of listing and transaction histories shapes ToM, seller adjustment, and execution outcomes in a tokenized real estate marketplace. The study contributes to the large real estate literature on the close relationship between pricing and ToM by using an innovative dataset with the complete sequence of listing price revisions for tokenized real estate traded on blockchain-based markets.

Data and methodology

The data for Paper 2 is built on the comprehensive secondary-market dataset of Paper 1, covering one of the most advanced real estate tokenization ecosystems of its time, RealT. More specifically, listing and transaction data from the most dominant market mechanism by trading volume, the P2P marketplace, is extracted and extended by reconstructing the full listing and selling paths of tokenized real estate on the marketplace. This complete listing and transaction dataset covers 111,180 listings, 72,917 sales, and 731 tokenized rental real estate properties in the United States and Latin America from 2020 to 2025. Methodologically, the paper employs standard OLS models and accelerated failure time (AFT) models on listing-level and revision-level data to study the determinants of ToM for tokenized real estate and the effect of price revision strategies on ToM and price outcomes when the full listing path is fully traceable for sellers and buyers.

Results and their contribution to science and practice

This paper is the first to extend the standard ToM literature and framework to tokenized real estate markets. While increasing narratives around real estate tokenization have emphasized its potential to alter market microstructure, as studied in Paper 1, or its implications for

investing and portfolio choice (see, e.g., Baum, 2021; Steininger, 2023), this study shows how the greater transparency enabled by tokenization affects market outcomes. In tokenized real estate markets, where the full listing history is public, both the magnitude and the path of price adjustment are first-order determinants of liquidity and execution outcomes. More specifically, ToM in this setting is strongly shaped by overpricing relative to a redemption benchmark (Buyback) provided by the tokenization platform. This confirms search-theoretic predictions that listings with higher and more persistent premia relative to the redemption benchmark remain longer on the market. Tokenization does not eliminate slow matching in real estate (Wheaton, 1990), but it links ToM more tightly to transparent relative pricing and observable platform liquidity. In other words, it more than negotiated property sales. A further key finding, one that is difficult to isolate in conventional real estate markets, is that, holding total price cuts constant, listings sold after a single decisive price cut achieve higher final premia than listings of the same underlying property repriced gradually. The results make three major contributions to science and practice. First, the paper provides the first large-scale empirical evidence on ToM, listing behavior, and seller adjustment for individual tokenized real estate listings, thereby shedding light on the market microstructure of fractionalized real estate traded on blockchain infrastructure. Second, this data-rich setting allows the study to re-examine classic questions of overpricing, marketing time, and seller adjustment in an environment where the posted price is the execution price and the full listing path is observable. These insights may also be relevant for traditional real estate listing platforms, where offer and price histories are becoming increasingly transparent to buyers. Third, the paper shows that the economic importance of tokenization lies not only in lot-size or liquidity transformation, but also in moving real estate exchange toward a more transparent, traceable microstructure in which relative pricing, market thickness, and publicly observable seller adjustments become more consequential for liquidity and execution outcomes.

Paper 3: Virtual Land in the Metaverse? Exploring the Dynamic Correlation with Physical Real Estate

Problems and objective

Lastly, Paper 3 provides a more futuristic glimpse into the potential of blockchain technology and tokenization for real-estate-like markets in purely virtual environments, often referred to as the metaverse. Blockchain-based metaverse platforms have attracted numerous companies and investors to purchase virtual land, represented by non-fungible tokens (NFTs), for commercial use, speculation, or status (Ante et al., 2023). Unlike earlier virtual worlds and their economies, in which ownership ultimately depended on centralized operators, blockchain-based virtual land can, in principle, reduce reliance on a single entity that can alter or shut down the virtual world, thereby reducing investor hold-up problems (Goldberg et al., 2024). More broadly, the relevance and research potential of virtual environments and purely digital yet spatial assets, as also reflected in economic work on internet domain names (Lindenthal, 2014; Lindenthal, 2018), continue to grow. Together with the notable parallels between virtual land and physical real estate, these developments motivate the research objective of this study, which is to investigate the interconnected dynamics between the two markets and shed light on their determinants and portfolio implications.

Data and methodology

The paper studies Decentraland, one of the earliest and largest blockchain-based metaverse platforms (Ordano et al., 2017), built on the Ethereum blockchain. The principal dataset comprises 15,855 secondary-market transactions of developed and undeveloped virtual land parcels, amounting to a trading volume of \$108 million between 2019 and 2023. Using this uniquely compiled dataset of virtual land sales, the paper constructs a repeat-sales index for virtual land, following Case and Shiller (1987), to measure its aggregate market returns. For the physical real estate market, the analysis relies on real estate investment trusts (REITs), specifically targeting commercial real estate due to the predominantly commercial setting of the observed virtual land market. To study the dynamic correlation between virtual and physical real estate market returns, the paper employs wavelet analysis, a time-frequency method frequently used in NFT-market settings (see, e.g., Dowling, 2022; Vidal-Tomás, 2022). To identify the key determinants of the correlation, the paper further estimates stepwise OLS regressions. Lastly, a portfolio analysis is performed to assess the risk-return implications of adding virtual land to real estate portfolios.

Results and their contribution to science and practice

The paper is the first to explore virtual land market returns in the metaverse and their correlation with physical real estate in the real world. To measure virtual land market returns, it adapts the repeat-sales index methodology from traditional real estate research, thereby making the performance of both markets empirically comparable. At a high level of abstraction, both assets serve as scarce spatial platforms within broader valuable networks, one in the digital realm and the other in the physical world (Goldberg et al., 2023). At the same time, the empirical results provide initial evidence that virtual land returns are largely independent from traditional real estate returns, suggesting that virtual land should not simply be understood as a virtual extension of the real estate asset market, but rather as a distinct real-estate-like virtual asset within the broader NFT investment class (Goetzmann et al., 2026). From a real estate perspective, one particularly interesting finding is that virtual land returns increased strongly during the Covid-19-induced lockdown period, in contrast to underperforming commercial real estate, resulting from restricted physical consumer activity, travel, and in-person interaction that partly shifted to online alternatives (Ghosh et al., 2023). The portfolio analysis indicates that this low level of connectedness suggests potential diversification benefits, especially for commercial real estate investors. However, these findings must be interpreted cautiously, given the short observation period. During periods of monetary tightening, by contrast, virtual land returns were highly volatile on the downside, and co-movement between virtual and physical real estate increased significantly, indicating that diversification benefits are time-varying and depend on the macroeconomic environment. Beyond these empirical insights, the study provides a hand-collected list and locational map of several dozen companies that acquired virtual land, offering early evidence on how firms experiment with commercial initiatives in the metaverse. Entrepreneurs and platform initiators may use those results to better understand virtual land market dynamics before launching new platforms or improving existing ones. Ultimately, the findings can help policymakers and regulators better categorize virtual land as a potential emerging asset class within the broader NFT and blockchain-based asset space.

5.2 Final Remarks and Outlook

Asset and capital markets are subject to constant efforts to improve efficiency and overall welfare. Within them, real estate markets, which are characterized by comparatively high levels of friction and inefficiency, therefore offer considerable potential for improvement. Over time, a variety of market innovations have been introduced, some of which have proven useful and become established, while others have remained more niche. What these innovations have in common is that they are becoming increasingly technological and fast-moving, often advancing more quickly than the financial, institutional, and regulatory structures surrounding them. This makes it ever more important for academia to remain at the forefront of these innovations and assess their economic, financial, and institutional implications. This cumulative dissertation responds to this need by shedding light on how blockchain and tokenization, among the most transformative technologies of recent decades, can reshape real estate markets.

Drawing on the open, though technically nontrivial, accessibility of public blockchain data, this dissertation aims to establish an empirically focused new field of research in blockchain-based real estate markets. Each study in this dissertation compiles a unique blockchain-based market dataset, comprising secondary-market panels for tokenized real estate, an extension of these panels with detailed listing and pricing histories, and virtual land sales in the metaverse. The datasets and their curation can serve as a cornerstone for applied blockchain-based market analytics, as public blockchain data becomes a novel and powerful source for real estate and asset research more broadly.

From an economic perspective, this dissertation contributes to the persistent debate on liquidity transformation for illiquid assets, particularly real estate, which has re-emerged in the context of blockchain and tokenization. The evidence presented in this dissertation suggests that tokenization facilitates secondary-market formation, but does not automatically create liquidity. Instead, liquidity remains conditional on sufficient trading demand, incentive-compatible liquidity provision, and the careful orchestration of coexisting execution mechanisms. That being said, the analysis of innovative blockchain-native mechanisms is important not because they solve liquidity transformation by design, but because their application to tokenized, formerly indivisible, heterogeneous, and slow-moving assets, such as real estate, reveals both their limits and potential areas for improvement. The findings highlight a central tension of tokenization as it aims to move slow-moving real assets into fast-moving, programmable trading environments, while the economic fundamentals of the underlying assets remain largely unchanged. A second major contribution is to show that the economic relevance of tokenization lies not only in liquidity, but also in transparency and traceability. Transparent, small-value token exchange moves real estate markets closer to financial-market microstructure, as relative pricing, market thickness, and revision histories become unusually observable and consequential for market outcomes. Finally, the analysis of virtual land reveals how real estate economic forces apply even when the physical asset disappears entirely. While real-estate-like market dynamics can emerge in purely virtual but spatial environments, these assets behave differently from physical real estate in terms of risk and return and should be understood as nascent assets within a broader blockchain-based virtual asset ecosystem. Taken together, blockchain does not make real estate frictionless, but rather changes how market frictions are represented, priced, observed, and empirically measured. This dissertation contributes to this understanding by examining real estate

market economics at the intersection of market microstructure research and public blockchain data.

As technological change comes with a high degree of uncertainty (Rosenberg, 1996), this uncertainty is also reflected in important limitations of this dissertation. One of its biggest limitations is therefore the early state of blockchain-based real estate markets. As a result, there is a limited availability of long observation periods, multiple comparable platforms, various public blockchain networks, and different real estate asset types. While the number of real estate tokenization ecosystems is increasing, token, platform, and blockchain standards remain highly fragmented and unevenly developed. With new tokenization data becoming available, future studies will be able to build on the approaches and findings of this dissertation to navigate this complexity more clearly. Furthermore, while this dissertation sheds light on the current realities of broader economic aspirations in real estate markets, especially liquidity and transparency, these dimensions do not automatically translate into welfare improvements. This dissertation is therefore limited in its ability to assess whether blockchain-based market mechanisms and their coexistence ultimately improve or harm welfare for market participants. Targeted welfare analyses would therefore help to derive sound policy recommendations on market interventions and regulation. Finally, although the dissertation focuses on real estate markets and aims to provide insights into tokenized illiquid assets more broadly, such generalizations should be made cautiously. Because comparable market data for other tokenized illiquid asset classes and private markets remain scarce, however, the insights can provide a valuable starting point for future research in this area.

The results of this dissertation can only provide an early glimpse into the outlook of blockchain and tokenization for real estate markets. Yet they already highlight a critical tension between technical feasibility and economic viability, pointing out that long-term adoption depends on more than new market technologies alone. However, creating welfare gains for market participants requires not only efficient market mechanisms but also legal and institutional frameworks that make the tokenized claims trustworthy and enforceable. The non-standardized nature of real estate creates important challenges for linking off-chain property realities to on-chain token markets. This opens several concrete future research avenues, the first of which concerns the design of reliable linking bridges, often referred to as oracles. Furthermore, the management intensity of real estate calls for research on blockchain-based governance designs to assign investors voting rights over collective property-management decisions. Beyond decentralized trading, which forms the heart of this dissertation, decentralized lending opens up another broad field for studying blockchain-based credit markets, especially for tokenized real estate, given the importance of leverage in traditional real estate finance. Taken together, the future research paths are vast and exciting and point to fundamental questions of how economic, legal, and institutional conditions could be altered through tokenization and blockchain-based market forms. Their exploration is essential to foster resilient markets that are liquid, trustworthy, and welfare-enhancing. As this dissertation shows, this aspiration has so far been only partially fulfilled in the context of tokenized real estate markets. While the growing institutional attention to blockchain and tokenization indicates that such potential exists, market selection will eventually take its course, separating the wheat from the chaff. Along the way, the task of academia will remain to help guide a path forward that is as stable and evidence-based as possible. Thanks to public blockchain technology, the empirical basis for this evidence is more accessible than ever.

References

- Ante L., Wazinski F.-P., and Saggu A. (2023). Digital real estate in the metaverse: An empirical analysis of retail investor motivations. *Finance Research Letters* 58, p. 104299.
- Barbon A. and Ranaldo A. (2026). On the quality of cryptocurrency markets: Centralized vs. decentralized exchanges. *Management Science. Articles in Advance*.
- Baum A. (2017). PropTech 3.0: The future of real estate. Report. Oxford: Saïd Business School, University of Oxford. URL: https://www.sbs.ox.ac.uk/sites/default/files/2019-03/FoRET-ReportSummary_0.pdf.
- Baum A. (2021). Tokenization – The future of real estate investment? *The Journal of Portfolio Management* 47.10, pp. 41–61.
- Bergkamp M., Sifat I., and Swinkels L. (2025). Market maturation and democratization effects of tokenized real estate. SSRN. Available at URL: <https://ssrn.com/abstract=5195172>.
- Buterin V. (2014). Ethereum: A next-generation smart contract and decentralized application platform. White Paper. URL: https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- Case K. E. and Shiller R. J. (1987). Prices of single-family homes since 1970: New indexes for four cities. *New England Economic Review* September/October, pp. 45–56.
- Castronova E. (2002). On virtual economies. SSRN. Available at URL: <https://ssrn.com/abstract=338500>.
- Cong L. W., Landsman W. R., Rabetti D., Zhang C., and Zhao W. (2025a). Tokenized stocks. SSRN. Available at URL: <https://ssrn.com/abstract=5937314>.
- Cong L. W., Mayer S., and Rabetti D. (2025b). Tokenizing real-world assets. Working Paper.
- Diamond D. W. and Dybvig P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91.3, pp. 401–419.
- Dowling M. (2022). Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters* 44, p. 102097.
- Geltner D. M., Miller N. G., Van de Minne A., Eichholtz P., Lindenthal T., and Shen L. (2026). *Commercial real estate analysis for investment, finance, and development*. 4th ed. Routledge.
- Ghosh I., Alfaro-Cortés E., Gámez M., and García N. (2023). Do travel uncertainty and invasion rhetoric spur metaverse financial assets? Gauging the role of media influence. *Finance Research Letters* 51, p. 103434.
- Goetzmann W. N., Huang D., and Nozari M. (2026). Non-fungible tokens as investment. Working Paper 34837. National Bureau of Economic Research.
- Goetzmann W. N., Spaenjers C., and Van Nieuwerburgh S. (2021). Real and private-value assets. *The Review of Financial Studies* 34.8, pp. 3497–3526.
- Goldberg M., Schär F., and Kugler P. (2024). Land valuation in the metaverse: Location matters. *Journal of Economic Geography* 24.1, pp. 67–92.
- Goldberg M., Schär F., and Thürkauf D. (2023). Retailing and customer engagement in the metaverse: An empirical analysis. SSRN. Available at URL: <https://ssrn.com/abstract=4547100>.

- Goldstein I., Jiang H., and Ng D. T. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics* 126.3, pp. 592–613.
- Gorton G. and Metrick A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics* 104.3, pp. 425–451.
- Harvey C. R., Rabetti D., Lin C., and Zhang C. (2025). Tokenized gold. SSRN. Available at URL: <https://ssrn.com/abstract=5995434>.
- International Monetary Fund (2026). Tokenized finance. IMF Note 2026/001. International Monetary Fund. URL: <https://www.imf.org/en/Publications/IMF-Notes/Issues/2026/04/10/tokenized-finance-565468>.
- Knight J. R. (2002). Listing price, time on market, and ultimate selling price: Causes and effects of listing price changes. *Real Estate Economics* 30.2, pp. 213–237.
- Kreppmeier J., Laschinger R., Steininger B. I., and Dorfleitner G. (2023). Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *Journal of Banking and Finance* 154, p. 106940.
- Lindenthal T. (2014). Valuable words: The price dynamics of internet domain names. *Journal of the Association for Information Science and Technology* 65.5, pp. 869–881.
- Lindenthal T. (2018). Monocentric cyberspace: The primary market for internet domain names. *The Journal of Real Estate Finance and Economics* 57.1, pp. 152–166.
- Liu A. and Chen C. (2025). From real estate financialization to decentralization: A comparative review of REITs and blockchain-based tokenization. *Geoforum* 159, p. 104193.
- Lukkarinen A. and Schwiendbacher A. (2023). Secondary market listings in equity crowdfunding: The missing link? *Research Policy* 52.1, p. 104648.
- Markowitz H. (1952). Portfolio selection. *The Journal of Finance* 7.1, pp. 77–91.
- Merlo A. and Ortalo-Magné F. (2004). Bargaining over residential real estate: Evidence from England. *Journal of Urban Economics* 56, pp. 192–216.
- Merton R. C. (1995). Financial innovation and the management and regulation of financial institutions. *Journal of Banking & Finance* 19.3–4, pp. 461–481.
- Nakamoto S. (2008). Bitcoin: A peer-to-peer electronic cash system. White Paper. URL: <https://bitcoin.org/bitcoin.pdf>.
- Ordano E., Meilich A., Jardi Y., and Araoz M. (2017). Decentraland: A blockchain-based virtual world. White Paper. URL: <https://decentraland.org/whitepaper.pdf>.
- Purnanandam A. (2011). Originate-to-distribute model and the subprime mortgage crisis. *The Review of Financial Studies* 24.6, pp. 1881–1915.
- Rosenberg N. (1996). Uncertainty and technological change. *Technology and Growth*. Conference Series 40. Boston, MA: Federal Reserve Bank of Boston, pp. 91–110.
- Savills (2025). How much is global real estate worth? Accessed 2 April 2026. URL: <https://impacts.savills.com/market-trends/how-much-is-global-real-estate-worth.html>.
- Schär F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review* 103.2, pp. 153–174.

- Schulte K.-W., Sebastian S., Allendorf G., Ropeter-Ahlers S.-E., and Leonhard H. (2026). Immobilieninvestition. Immobilienökonomie I: Betriebswirtschaftliche Grundlagen. Ed. by K.-W. Schulte, S. Bone-Winkel, and W. Schäfers. De Gruyter Oldenbourg, pp. 725–838.
- Smith A. (1776). An inquiry into the nature and causes of the wealth of nations. Vol. 1. London: W. Strahan and T. Cadell.
- Steininger B. I. (2023). Return–risk analysis of real estate tokens: An asset class of its own. *The Journal of Portfolio Management* 49.10, pp. 83–102.
- Swinkels L. (2023). Empirical evidence on the ownership and liquidity of real estate tokens. *Financial Innovation* 9, pp. 1–29.
- Vidal-Tomás D. (2022). The new crypto niche: NFTs, play-to-earn, and metaverse tokens. *Finance Research Letters* 47, p. 102742.
- Wheaton W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy* 98.6, pp. 1270–1292.
- Whitaker A. and Kräussl R. (2020). Fractional equity, blockchain, and the future of creative work. *Management Science* 66.10, pp. 4594–4611.
- Williams J. T. (1995). Pricing real assets with costly search. *The Review of Financial Studies* 8.1, pp. 55–90.
- Wood G. (2014). Ethereum: A secure decentralised generalised transaction ledger. Yellow Paper. URL: <https://ethereum.github.io/yellowpaper/paper.pdf>.