



Ph.D. Thesis

Decentralized Finance

Building and Analyzing Financial Infrastructure on Blockchain Technology

Carl Victor von Wachter

Advisor: Prof. Dr. Omri Ross

Submitted: 31.07.2023

This thesis has been submitted to the Ph.D. School of the Faculty of Science, University of Copenhagen

Abstract

Decentralized Finance (DeFi) has experienced astonishing surges in user adoption and asset accumulation. The financial assets secured in the broader DeFi ecosystem grew from \$600 million in January 2020 to \$38 billion in January 2023 - an increase of no less than 6,333% over three years. *Decentralized Finance* refers to an ecosystem of financial applications on public blockchain technology that offers financial services and instruments while reducing the dependency on intermediaries. DeFi's overarching aim is to rethink traditional financial infrastructure by minimizing the power of single entities as well as providing more accessible and efficient financial services.

Despite the promising potentials and genuine objectives, building a financial system on a blockchain remains challenging owing to the technology's infancy and inherent technical compromises. DeFi also presents a distinct environment for financial services, one that is openly accessible, continually operating, and adversarial. These conditions may have significant and distinct impacts on the broad range of financial services, further characterized by regional differences.

This cumulative thesis consists of an essay that investigates the extent to which blockchain technology is suitable for the implementation of a financial service as well as six publications. It consolidates the contributions of these six publications, each developed during a three-year Ph.D. program, into an assessment framework that proposes guiding principles that determine when the use of blockchain technology is advantageous for a financial service. Categorized by the overarching domain of DeFi, the individual publications follow a research approach based on prototyping or network analysis. The built prototypes serve as proofs of concept and aid in understanding the problem. The analytical work explores the de facto use of existing applications, aiming to broaden the discussion from purely technical to encompass human interactions and economics. Specifically, the publications present blockchain-based building blocks for standard financial services such as targeted stimulus payments, asset reserves, equity, and crowdfunding. The empirical publications propose algorithms to quantify domain-specific phenomena such as system integration, illicit behaviors in marketplaces, and financial arbitrage.

This thesis contributes to the existing knowledge, particularly to Information Systems (IS) research, by proposing an evaluation framework and increasing the understanding of blockchain technology's applicability to financial services. Further, the individual publications contribute to their respective field by presenting applicable blockchain-based building blocks for financial services and domain-specific algorithms for the DeFi ecosystem.

With the impressive growth of DeFi, amassing billions of USDs in assets, the significance of this research continues to evolve, potentially addressing critical questions toward establishing a more efficient and resilient financial system based on blockchain technology.

Resumé

DeFi har set betydelig stigning i brugeradoption og vækst i investerede aktiver. De finansielle aktiver sikret i det bredere DeFi økosystem voksede fra \$600 millioner i januar 2020 til \$38 milliarder i januar 2023 - en stigning på ikke mindre end 6.333% over tre år. Terminologien *Decentralized Finance* refererer til et økosystem af finansielle applikationer på offentlig blockchain teknologi, der tilbyder finansielle tjenester og instrumenter, samtidig med at afhængigheden på *mellemmænd* reduceres. Det overordnede mål bag DeFi bevægelsen er at minimere enkeltentiteters magt og samtidig tilbyde mere tilgængelige og effektive finansielle tjenester. Til trods for det lovende potentiale, er det stadig udfordrende at bygge et finansielt system på en blockchain på grund af teknologiens ungdom hvilket introducerer tekniske kompromiser. DeFi præsenterer også et særpræget miljø for finansielle tjenester, et der er åbent tilgængeligt, konstant i drift og modstanderorienteret. Disse forhold kan have betydelige og unikke virkninger på det brede udvalg af finansielle tjenester, yderligere karakteriseret ved regionale forskelle.

Denne kumulative afhandling består af et essay, der undersøger i hvilket omfang blockchain teknologi er egnet til implementering af finansielle tjenester, samt seks publikationer. Afhandlingen sammenfatter bidragene fra disse seks publikationer, hver udviklet i løbet af en treårig Ph. D. proces. Foreslår vejledende principper, der afgør, hvornår brugen af blockchain teknologi er fordelagtig for en finansiell tjeneste. De individuelle publikationer følger en forskningstilgang baseret på prototyping eller netværksanalyse, under det overordnede tema, DeFi. Prototyperne fungerer som proof-of-concept. Det analytiske arbejde udforsker brugen af eksisterende applikationer med formålet at udvide diskussionen til at omfatte menneskelige interaktioner med teknologien. Specifikt præsenterer publikationerne blockchain-baserede byggeklodser for finansielle programmer såsom målrettede stimulusbetalinger, egenkapital og crowdfunding. De empiriske publikationer foreslår algoritmer til at kvantificere domænespecifikke fænomener såsom systemintegration, ulovlige opførsel på markedspladser og finansiell arbitrage.

Denne afhandling bidrager til den eksisterende viden ved at foreslå et destilleret rammearbejde der øger forståelsen af blockchain-teknologiens anvendelighed i finansielle tjenester. Yderligere bidrager de individuelle publikationer til deres respektive felt ved at præsentere anvendelige byggesten og domænespecifikke algoritmer for DeFi-økosystemet.

Med den imponerende vækst af DeFi, der akkumulerer milliarder af dollars i aktiver, fortsætter denne forsknings betydning med at udvikle sig, og kan potentielt adressere kritiske spørgsmål ved at etablere et mere robust finansielt system baseret på blockchain-teknologi.

Preface

This thesis is a consolidation of six individual contributions, each developed throughout my Ph.D. program at the University of Copenhagen's Department of Computer Science (DIKU) from August 2020 to July 2023. The Ph.D. program coincided with the surge of DeFi's popularity since 2020, attracting fascinating minds, fostering global cooperation, and offering intriguing research opportunities. The research was driven by my passion for building and analyzing DeFi applications to increase my understanding of blockchain-based finance.

This thesis's primary objective is to consolidate the findings of the individual publications into a comprehensive assessment framework that seeks to establish guiding principles that determine when the use of blockchain technology is advantageous for a financial service. Given the inconclusiveness when implementing a financial service on blockchain technology, the framework aims for a balanced evaluation, drawing generalizable insights from the publications and providing a perspective with relevant criticism. While collectively focusing on DeFi, two main methodologies categorize the selected publications: prototyping and network analysis. The built prototypes serve as proofs of concept and aid in understanding the problem. The analytical work retrieves the public blockchain data and explores the de facto use of existing applications, so as to broaden the discussion from purely technical to encompass human interactions and economics. Thus, the research reflects the trajectory of my Ph.D. journey from a technical perspective to an analytical one. Further, it is representative of the interdisciplinary aspects of blockchain technology and DeFi.

This thesis has two parts. Part 1 is an essay that consolidates the findings from the six individual publications; it is organized in five chapters. Chapter 1 serves as an introduction to the topic and presents the research question. Chapter 2 lays the technical foundations of blockchain technology and DeFi, further summarizing the technical properties of blockchain technology and DeFi's characteristics. Chapter 3 describes this thesis's research approach, discusses the publications and their contribution to the thesis, and introduces the applied research methodologies. Chapter 4 starts with deriving the evaluation items for the framework from the individual publications. Subsequently, the chapter presents the complete framework and conducts a brief evaluation. Lastly, the framework and DeFi is critically discussed. Chapter 5 concludes and suggests further research avenues. Part 2 of this thesis compiles the six publications, including five peer-reviewed papers and one currently under review. This article-based thesis presents the papers in their published format, with the exception of formatting adjustments and the correction of typos. The following publications, sorted by peer-reviewed publication date, are included in this thesis:

1. **An Introduction to Decentralized Finance (DeFi)**
Johannes Rude Jensen, Victor von Wachter, Omri Ross
Complex Systems Informatics and Modeling Quarterly (CSIMQ), 2021.
2. **Measuring Asset Composability as a Proxy for DeFi Integration**
Victor von Wachter, Johannes Rude Jensen, Omri Ross
International Conference on Financial Cryptography and Data Security (FC), 2021. International Workshops.
3. **NFT Wash Trading - Quantifying Suspicious Behavior in NFT Markets**
Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, Omri Ross
International Conference on Financial Cryptography and Data Security (FC), 2022. International Workshops.
4. **Blockchain-based Infrastructure for Emerging Economies**
Johannes Rude Jensen, Victor von Wachter, Omri Ross
European Conference on Information Systems (ECIS), 2022.
5. **Kickstarting Blockchain: Designing Blockchain-based Tokens for Equity Crowdfunding**
Tobias Guggenberger, Benjamin Schellinger, Victor von Wachter, Nils Urbach
Electronic Commerce Research (EC), 2023.
6. **Fundamentals of Perpetual Futures**
Songrun He, Asaf Manela, Omri Ross, Victor von Wachter
This paper is currently under peer review.

Acknowledgements

My Ph.D. journey would not have been possible without the support and encouragement of many remarkable individuals, to whom I owe a great deal of gratitude. I am indebted to many people who supported my academic progress during my Ph.D.

I am extremely grateful to my supervisor Prof. Omri Ross. Your unwavering enthusiasm for DeFi has been a source of inspiration and motivation. You gave me the freedom and confidence to pursue my passion and have played a crucial role in my development as a researcher, teacher, and individual.

I am very thankful to Prof. Asaf Manela for your warm hospitality and academic supervision during my research project in Tel Aviv, Israel. Your guidance exposed me to a whole new perspective on academic research. It was also an encouraging opportunity to experience and appreciate a different culture.

I am indebted to my fellow researcher, Johannes. You were an outstanding part of my journey. Your companionship, support, and countless hours of intellectually stimulating discussions have enriched my experience. Without you, this journey would have been much more challenging.

My friend Ferdinand, your influence introduced me to the intriguing world of blockchain technology and you also planted the idea to pursue a Ph.D. In many research ideas, your constructive and technical feedback has been instrumental.

I also wish to express my gratitude to my colleagues at PLTC, particularly Henrik, Fritz, Cosmin, Martin, and Ken. You supported me throughout the journey and shared many fascinating discussions, excursions, and conversations.

I am grateful for receiving a scholarship from the European Union's Horizon 2020 research and innovation program, under the Marie Skłodowska-Curie grant. It led to my relocation from Munich, Germany to Copenhagen, Denmark, which brought its own unique experiences. Navigating the transition to remote research, teaching, and socializing amidst the challenging times of the COVID-19 pandemic was not easy. I am deeply appreciative of the personal support received during this challenging time.

My parents are always there for me and deeply believe in my abilities. Your unwavering love and support are the foundation of who I am today.

My sisters, with their patience, love, and cheerful children are my sanctuary in stressful times.

And finally, I am grateful to Annika for your unwavering love, your endurance for living in two different countries, the shared passion for exotic countries, and even your tolerance to laugh at my lame jokes. I cannot emphasize enough how your presence made this journey less strenuous.

Table of Contents

Abstract	i
Resumé	ii
Preface	iii
Acknowledgements	v
Table of Contents	vi
List of Tables	viii
List of Figures	ix
1 Introduction	1
2 Foundation	4
2.1 Blockchain Technology	4
2.2 Decentralized Finance	8
2.3 Properties of Blockchain Technology and Characteristics of DeFi	13
3 Research Approach of the Thesis	21
3.1 Publications	22
3.2 Methodologies	27
4 Toward an Assessment Framework	30
4.1 Items from Building DeFi Applications	30
4.2 Items from Analyzing DeFi Applications	38
4.3 The Assessment Framework	42
4.4 Evaluation	46
4.5 Discussion	50
5 Conclusion	57
References	60
Paper Overview	73
A Introduction to Decentralized Finance	77

<i>TABLE OF CONTENTS</i>	vii
Paper References	87
B Measuring Asset Composability for DeFi Integration	90
Paper References	96
C NFT Wash Trading	98
Paper References	110
D Blockchain-based Infrastructure for Emerging Economies	113
Paper References	124
E Kickstarting Blockchain: Blockchain-based Equity Token	128
Paper References	160
F Fundamentals of Perpetual Futures	168
Paper References	209

List of Tables

2.1	Properties of Blockchain Technology	14
2.2	Characteristics and Values of DeFi	15
4.1	The Assessment Framework	45
4.2	Framework Evaluation	49
A.1	DeFi Agent Classification	81
B.1	Transactions of DeFi Assets during 2020	93
C.1	Overview of the Results	103
C.2	Results for each NFT Collection	108
D.1	The Artifact Requirements	117
D.2	Artifact Evaluation	122
E.1	ERC Token Standards on Ethereum	133
E.2	Areas for Improvement	139
E.3	The Design Objectives	143
E.4	Criteria-based Evaluation of Affs and DOs	151
E.5	Overview of the Interviewees	152
E.6	Design Principles for Blockchain-based Equity Tokens	157
F.1	Discounted Payoffs to Arbitrage Strategies	177
F.2	Sample Descriptions	180
F.3	Trading Costs Specifications	182
F.4	Performance of Random-maturity Arbitrage Strategy	187
F.5	Performance Over Time: High Trading Costs Tier	188
F.6	Return Decomposition: Price Convergence vs Funding Rate Payment	189
F.7	Regression of the Futures-Spot Gap	191
F.8	Portfolio Performance: High Trading Cost	204
F.9	Portfolio Performance under Different Trading Costs	205
F.10	Return Decomposition: Price Convergence vs Funding Rate Payment	206

List of Figures

2.1	Blocks and Transactions of a Blockchain	5
2.2	A Layered Model of DeFi	10
3.1	The Research Approach of the Essay	23
3.2	A schematic DSR Approach	28
A.1	DeFi Applications on Permissionless Blockchain	80
A.2	AMM Price Discovery	82
B.1	Method and Asset Tree Structure	92
B.2	Financial Integration of Assets	94
C.1	Closed Cycle Pattern	102
C.2	Rapid Sequence no Market Risk Pattern	102
C.3	Elapsed Time for Closed Cycles	104
C.4	Wash Trading with Respect to Lifetime	107
C.5	Trades and unique Trade Partners	109
D.1	The cyclical DSR Workflow	116
D.2	Artifact Overview	118
D.3	Network Analysis of Transactions	120
E.1	Research Process	136
E.2	Concretization and Abstraction	137
E.3	Class Diagram	144
E.4	Sequence Diagram Issuance	147
E.5	Sequence Diagram Transaction	147
E.6	Iterative Design, Development, and Evaluation of the Artifact	153
E.7	Service Ecosystem	156
F.1	Total Trading Volumes of Perpetual Futures	173
F.2	Bitcoin annualized Funding Rate	174
F.3	Deviations of Perpetual Futures from no-arbitrage	183
F.4	Correlation of ρ	184
F.5	Random-maturity Arbitrage Strategy: Bitcoin	186
F.6	Trading View of Perpetuals on Binance	193

F.7	Trading Cost Tiers Spot	194
F.8	Trading Cost Tiers Perpetuals	195
F.9	Ether annualized Funding Rate	196
F.10	Annualized Interest Rate from Aave	197
F.11	Trading Strategy Visualization: Random-maturity Arbitrage	198
F.12	Trading Strategy Visualization: No Trading Costs	199
F.13	Trading Strategy Visualization: Low Trading Costs	200
F.14	Trading Strategy Visualization: Medium Trading Costs	201
F.15	Trading Strategy Visualization: High Trading Costs	202
F.16	Two-threshold Trading Strategy: Bitcoin	203
F.17	High-frequency Even Study around the Funding Rate Payment	208

Acronyms

AML anti-money laundering

AMM automated market maker

BTC Bitcoin

DeFi Decentralized Finance

DSR Design Science Research

EoA externally owned address

ERC Ethereum request for comment

ETH Ether

IS Information Systems

NFT nonfungible token

NGO non-governmental organization

TpS transactions per second

USD United States Dollar

VM virtual machine

Introduction

Since 2020, DeFi - an ecosystem of applications catering financial services on public blockchain technology [106, 51, 124] - has experienced astonishing increases in user adoption and asset accumulation. This growth is best exemplified by the expansion of financial assets secured in the broader DeFi ecosystem, which grew from \$600 million in January 2020 to \$38 billion in January 2023 - an increase of no less than 6,333% over three years¹. DeFi uses public blockchain technology with smart contract capabilities to facilitate financial services through autonomous programs, thereby reducing the dependency on financial intermediaries [106, 43]. These autonomous programs are typically open-source. Public access for inspection, copying, and modification, ultimately seeks to establish trust in the program's integrity [106, 5]. Another defining characteristic is DeFi's composability [65, 47, 8]. Owing to the underlying blockchain technology, financial services can integrate with one another leading to a single, coherent financial system [109]. Users benefit from this composability since they can seamlessly move their assets between different applications, thereby finding customized financial services that cater to their preferences [Paper 2: 121]. This results in a simultaneously competitive and collaborative financial system that optimizes the distribution of assets under strong market forces [51]. DeFi's overarching aim is to rethink traditional financial infrastructure by minimizing the power of single entities and a focus on more accessible, transparent, and efficient financial services [51, 106, 44, 23].

Despite its potentials and genuine objectives, developing financial applications for the DeFi ecosystem remains challenging owing to the nascent state of the technology and its inherent technical trade-offs [133, 132]. Further, DeFi presents a distinct environment for financial services, one that is openly accessible, continually operational, and adversarial [23, 25]. These conditions may have significant and distinct impacts on the broad range of financial services, further characterized by regional differences [28, 7]. Thus, while blockchain technology's applicability varies across different financial services, each financial service may benefit from individual assessment. Despite these techno-

¹<https://defillama.com>, accessed 16th May 2023

logical nuances, the uniqueness of individual financial services, and regional differences, this essay hypothesizes that there are guiding principles that determine when the use of blockchain technology is advantageous for a financial service. Thus, this essay asks: To what extent is blockchain technology suitable for the implementation of financial services?

This essay addresses this research question by consolidating the individual contributions of six publications, each developed during my Ph.D. process. These publications, while collectively focusing on DeFi, can be categorized based on their research approach of prototyping or network analysis. The prototypes built in the publications serve as proofs of concept and aid in understanding the problem at hand [56]. The analytical papers explore the de facto use of existing applications, aiming to broaden the discussion from purely technical to encompass human interactions and economics [119, 86]. Specifically, the publications present blockchain-based building blocks for standard financial services in the areas of targeted stimulus payments, asset reserves, equity, and crowdfunding. The empirical publications propose algorithms, quantifying domain-specific phenomena such as system integration, illicit behaviors in marketplaces, and arbitrage.

This essay culminates in an assessment framework that proposes guiding principles that determine when the use of blockchain technology is advantageous for a financial service. The framework consists of 24 evaluation items drawn from theoretical and methodological arguments of the individual publications. With a scoring-based approach, it seeks to provide a balanced evaluation that can serve as a valuable tool for both researchers and practitioners, identifying key questions, relevant literature, and remaining challenges. The proposed assessment framework seeks to contribute to the existing knowledge, especially in the IS discipline - an interdisciplinary field that combines elements of computer science, economics, and social science [9, 105]. As blockchain technology is a combination of distributed systems, cryptography, and economic incentives [38, 133, 70], scholars argue that it is well-suited for an interdisciplinary research approach [112, 90]. Existing assessment frameworks have explored blockchain technology generically [132, 58, 94, 130, 78, 42, 93, 69, 71]. [69] calls for the expansion of blockchain-based frameworks to include domain-specific questions, particularly highlighting economic aspects. Further highlighting the impact of blockchain technology on financial services, [112, 82, 43] motivate to explore how financial applications can be designed to capitalize on the strength of blockchain technology. Thus, there remains a gap in the literature for a framework that is specifically tailored for financial services on blockchain technology. With a focus on financial services, the presented framework is a structured approach to objectify blockchain technology's applicability to financial services. Further, each of the six publications seeks to contribute to its respective field by proposing blockchain building blocks for financial services or algorithms to quantify domain-specific phenomena in the DeFi ecosystem.

With the impressive growth of DeFi, amassing billions of United States Dollar (USD)s in assets, the significance of this research continues to evolve, potentially addressing critical questions toward establishing a more resilient financial system based on blockchain technology.

The focus of my Ph.D. program and thus this essay is the exploration of financial services on public blockchain technology from a technical and an analytical perspective. The Ph.D. program coincided with the surge of DeFi's popularity since 2020. Popular public blockchains such as Bitcoin [87] and Ethereum [19, 129] have established strong network effects, with a large number of developers, users, and applications that continually experiment with the technology [106, 23, 54]. This public financial ecosystem offers fascinating and dynamic research opportunities and sets limitations to the scope. Permissioned blockchain networks have not been a focal point of the thesis [97, 54]. Further, financial services are also subject to a variety of regulatory considerations. Regulatory aspects are only added in some parts for comprehensiveness. For an extensive analysis, readers are referred to the works of [135, 91, 16, 26].

This thesis has two parts. Part 1 is an essay that consolidates the findings from the six individual publications; it is organized in five chapters. Chapter 1 serves as an introduction to the topic and presents the research question. Chapter 2 lays the technical foundations of blockchain technology and DeFi, further summarizing the technical properties of blockchain technology and DeFi's characteristics. Chapter 3 describes this thesis's research approach, discusses the publications and their contribution to the thesis, and introduces the applied research methodologies. Chapter 4 starts with deriving the evaluation items for the framework from the individual publications. Subsequently, the chapter presents the complete framework and conducts a brief evaluation. Lastly, the framework and DeFi is critically discussed. Chapter 5 concludes and suggests further research avenues. Part 2 of this thesis compiles the six publications, including five peer-reviewed papers and one currently under review. This article-based thesis presents the papers in their published format, with the exception of formatting adjustments and the correction of typos.

Foundation

Blockchain technology has captured considerable global attention over the past decade [70, 114, 138]. First introduced in 2008 as the underlying technology for Bitcoin [87], a digital currency, blockchain has since evolved into a versatile technology for applications spanning various industries.

This section provides an overview over the technical fundamentals of blockchains, drawing on an extensive body of literature [5, 4, 133, 70, 138]. Various blockchains exist and discussions of the technical nuances can become very complex. This section utilizes frameworks and taxonomies in the IS research field that helps to abstract and structure this emerging technology [72, 132, 40, 111].

2.1 Blockchain Technology

A blockchain is a distributed ledger in which a large network of untrusted participants synchronizes and stores a shared state following a strict protocol [31, 39]. Each participant maintains a replica of the blockchain database [133, 70]. The participants are financially incentivized to operate the system through the issuance of a native protocol asset, such as Bitcoin (BTC) on the Bitcoin blockchain and Ether (ETH) on the Ethereum blockchain. Typically, a blockchain is used to account for the transfers of these native protocol assets, essentially creating a digital currency [114]. Signed messages, or *transactions*, are integral to any blockchain. Transactions sent to the network represent state transitions and are bundled into *blocks* [4, 5, 70]. Each block contains multiple transactions and a hashed representation of the previous block. The blocks are cryptographically linked and in chronological order. To alter transactions in a previous block, without invalidating the chain, is prohibitively expensive since it requires changing the full chain of hashes. Thus, transactions are never altered; instead, new transactions are added to the blockchain, creating an append-only database [132, 70, 133]. Figure 2.1 schematically depicts a blockchain.

Blockchain technology is a decentralized computer system. [55, 133] distinguished between technical and organizational decentralization. Technically, a blockchain is a distributed system hosted on servers by a large number of

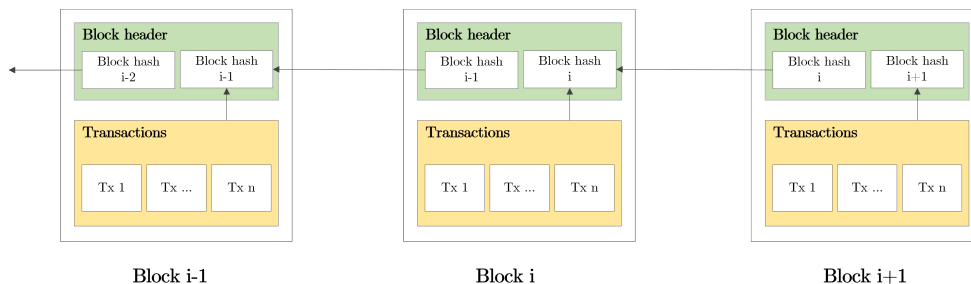


Figure 2.1: The blocks and transactions of a blockchain. The blocks are cryptographically linked with a reference to the previous block hash. Each block stores multiple transactions. This creates an immutable database where the transactions are stored in chronological order.

participants, but it is also peer-to-peer, as these participants are independent entities [55]. The system is organizationally decentralized as no single participant has privileged access. The system is governed for software updates or conflict resolutions by the participants who operate and use the system. This stands in contrast to common server-hosted systems, where the server host has highly privileged access [55]. Network participants are financially incentivized to store, propagate, and verify blocks. When one participant announces a block with a list of transactions, the majority of the other network participants must verify and confirm the new block. This process of reaching an agreement on state changes is referred to as *consensus*. In broader terms, consensus rules ensure that participants in a distributed system converge on the same system-wide state [137, 70, 40].

Consensus protocols are sophisticated game-theoretical mechanisms that promote positive behaviors and penalize hostile ones in an open environment. These consensus protocols are generally designed to be tolerant of Byzantine behaviors [73], which assumes that peers can become unavailable, adversarial, or communicate contradictory information to other peers. Another related concept is *finality*, which provides assurance that a transaction on the blockchain cannot be reversed [3, 70]. A committed block is not immediately immutable; altering a transaction requires changing and propagating all subsequent blocks [132, 133]. Thus, the older a block is, the more likely it is to be considered final under *probabilistic finality* [70], Bitcoin's finality concept. In practice, a few blocks are typically sufficient to provide economic guarantees that a block will not be reversed. The ability to achieve consensus in a distributed network under hostile conditions without centralizing control is the core principle of all public blockchains [5, 133]. Participants do not have to trust one another, only the blockchain protocol rules. Similar in principle, the exact rules for consensus and finality differ among the blockchain protocols

[132, 137, 40]. The specifications of each blockchain protocol are implemented in open-source clients. Updating the rules in the form of a new software version requires agreement from the majority of the network participants [4, 5].

The block size and block time are fundamental to a blockchain's throughput [24], with block sizes typically in the megabytes and block times ranging from sub-seconds to a few minutes [5]. In anticipation of high demand, blockchains seek to process large numbers of transactions quickly, potentially handling thousands of transactions per second (TpS). However, there are technical barriers to scaling blockchains [24, 77], since each block must be propagated and agreed upon in the network. The majority of peers must receive, process, store, and send verification for each block. To use the network, users must create an account. Blockchains utilize public key cryptography to identify participants and authorize asset ownership. After generating a public and private key pair, the account's unique identifier, an *address*, can be derived by applying a hash function to the public key. The corresponding private key is used to create a digital signature, which directly controls access to the account's assets. Ubiquitous software *wallets* helps the users to sign and broadcast transactions, abstracting away cryptographic complexity [4, 5, 70].

Access to a blockchain can be either permissioned or permissionless. Permissionless blockchains, such as Bitcoin and Ethereum, offer a large number of participants unrestricted and unauthenticated access. Anyone can read and write transactions [123, 71]. In practice, the openness of permissionless blockchains has led to widespread distribution and rapid innovation, resulting in recent blockchain trends being primarily associated with permissionless networks [71]. Various instantiations of blockchain technology exist. Blockchains follow socially agreed-upon philosophies and objectives. For instance, blockchains may prioritize performance, security, or decentralization [132, 40]. The various instantiations of blockchain technology undergo rapid and ongoing development.

In 2013, [19] proposed adding a virtual machine (VM) to the shared state of a blockchain. A VM is an execution environment that computes and updates valid, general-purpose state changes, enabling the execution of arbitrary complex programs [5, 70]. [129] formalized the proposed concept for a trust-minimized execution environment and database, which was subsequently implemented as Ethereum. Network participants pass instructions to the VM using a higher-level programming language designed for writing programs [5]. A computer program executed on the VM is often called a *smart contract*, a term coined by [89]. Similar to classes in object-oriented programming, smart contracts are containers that include state variables as persistent data and functions to manipulate these [111, 12]. Thus, blockchains with smart contract capabilities can be conceptualized as transaction-based state machines, containing millions of executables, each with its own conditional logic and permanent storage [5, 70]. Although smart contracts are potentially less efficient compared to centralized computing, their advantage lies in protocol-controlled,

deterministic execution and storage. The logic can be verified, and it runs exactly as expected, in contrast to regular server-based web applications, which require user trust [133, 106, 12]. Once deployed the smart contract is shared across the network, allowing anyone to interact with it. While storage updates occur when the smart contract is called, the implemented logic cannot be altered - only new instances can be deployed. A smart contract is also identified by a unique address. In contrast to the externally owned address (EoA), a smart contract account only contains the program and lacks the ability to autonomously initiate transactions [133]. Smart contracts are ultimately executed as a result of a transaction invoked by EoAs, reacting solely to the received instructions. When an external participant sends a transaction to a contract address, it prompts the contract to run in the VM [5, 70]. Notably, contracts can call other contracts, generating arbitrary execution paths of varying depths. These cascades of internal transactions, also known as traces, form the technical foundation for the composability of decentralized applications [65].

To maintain a targeted block size, blockchain VMs are constrained by a finite number of computational steps for all transactions within a block. Since all participants must replicate computations for validity, computational steps are limited and priced [70, 5]. The cost of each operation within a transaction is metered and paid upfront, minimizing adversarial activities and excessive usage. Participants compete for these limited resources, resulting in auction-style pricing of computation and storage based on demand. The fees are higher in times of high demand and lower in times of low demand [30, 25]. VMs execute computational steps atomically, i.e. if conditions in any called contract are not met, execution halts and the transaction is abandoned. In the case of failed execution, the transaction is still recorded, but all state changes are rolled back, and the computation fee is deducted from the originating account [65, 70].

A blockchain, as a data structure, is a public, highly redundant, immutable list of transactions within a sequence of blocks, sorted by timestamp in chronological order [133, 70]. Owing to its characteristics, blockchains provide a unique dataset [109, 119]. While the block-based design is generic, the exact data structures and encoding are blockchain-specific and typically highly optimized. As the blockchain is replicated across all participants and encompasses millions of transactions, data storage is a valuable resource [70, 5]. Storing information is a basic operation in the VM and as such is also subject to block limits and fees. Applications typically seek to reduce or outsource a transaction's data footprint. Large data, such as images, are often stored *off-chain*, i.e. outside of the blockchain environment [5]. Conversely, some applications require external information [39]. The execution environment of a blockchain is endogenous, with VMs only accessing the limited information present in current or past blocks [132]. External data, such as stock prices or currency conversion rates, are not accessible and must be fed into the blockchain. *Ora-*

cles are specialized smart contracts that provide relevant external information to the blockchain environment for other applications to access [2, 32].

Since access to the blockchain is unrestricted, anyone can view the participants' balances. To maintain privacy, public key cryptography is used as an anonymization technique. Without voluntary disclosure, the true identity behind an address remains unknown [70, 5, 4, 133]. Further, as participants can create an arbitrary number of addresses, it is not possible to determine whether a participant holds multiple addresses. Thus, data are transparent but pseudonymous. This presents an intriguing paradox: pseudonymous identities protect privacy, while simultaneously making it hard to prevent the obfuscation of illicit practices [Paper 3: 122].

Introduced in 2008, blockchain technology has received interest from both practitioners and academia owing to its intriguing combination of technologies. Blockchains represent a blend of various technologies within the computer science discipline, including consensus, public key cryptography, data structures, and hashing [70, 38, 133]. A smart contract-enabled blockchain contains a computation engine to process general-purpose transactions. Diverse design philosophies have given rise to a wide array of blockchains, which continue to undergo rapid development. Notably, public blockchains with smart contract capabilities have attracted significant attention, fostering a thriving ecosystem of applications. The following section will concentrate on these applications built on smart contract blockchains.

2.2 Decentralized Finance

Enabled by their technical characteristics, blockchains have attracted significant interest in a variety of use cases. The financial industry is particularly receptive to blockchain technology, given the importance of trust and security [74, 113, 51, 23].

Blockchains evolve beyond simply facilitating the transfer of native digital currency. Smart contract blockchains emerged as a platform for programmable decentralized applications establishing a thriving ecosystem [19, 124]. DeFi generally refers to an ecosystem of applications built on permissionless blockchain technology, specifically catering to financial services [106, 135, 124]. Although various blockchains exist, DeFi is typically associated with Ethereum, since it has the largest ecosystem in terms of users, available applications, and development activities [124, 106][Paper 1: 61]. These applications typically replicate common financial services, such as exchanging assets, borrowing and lending, insurance, and derivatives [48, 124, 106][Paper 1: 61]. Further, blockchain technology's unique properties have led to the invention of novel financial services [Paper 1: 61], such as automated market maker (AMM) [131, 13] and flash loans [99, 96]. DeFi uses public blockchain technology with smart contract capabilities to facilitate financial services through

autonomous programs and thereby reducing the dependency on financial intermediaries [106, 43]. DeFi's overarching aim is to rethink traditional financial infrastructure by minimizing the power of single entities and a focus on more accessible, transparent, and cost-effective financial services [51, 106, 44, 23].

DeFi's rising popularity can be measured using blockchain data, with key metrics including the total economic value secured in DeFi applications [124, 106]. This growth is best exemplified by the expansion of financial assets secured in the broader DeFi ecosystem, which grew from \$600 million in January 2020 to \$38 billion in January 2023 - an increase of no less than 6,333% over three years¹. Notably, in a seminal paper [109] positions DeFi as the largest public dataset for financial markets, owing to the transparency and detail inherent in blockchains. It enables granular detail on individual trades and users - a fascinating opportunity for academic research. This transparency level contrasts with traditional financial infrastructure, where the acquisition of granular empirical data is increasingly challenging owing to regulatory constraints or intellectual property rights [109].

Conceptually, DeFi can be envisioned as a multi-layered architecture [106, 46, 8]. Informed by [Paper 1: 61], Figure 2.2 depicts the layered understanding of DeFi. The blockchain is the lowest layer and serves as a settlement layer for state-changing transactions powered by its native protocol asset. The blockchain ensures that any state changes adhere to the protocol rules and fundamentally facilitates confidence in the computational system [27]. The upper layers are empowered by a VM on top of a blockchain, facilitating applications that require general-purpose transactions. In the protocol layer, smart contracts are employed to create basic financial objects in the form of digital tokens [104, 106][Paper 5: 49]. These standardized financial objects are often shared across the DeFi ecosystem. Decentralized applications utilize these standardized shared assets to create sophisticated financial services. These applications are implemented through a set of smart contracts in the application layer. The smart contracts implement application-specific rules ultimately offering a financial service [106, 46]. These applications can be structured in verticals, such as trading, lending, insurance, and derivatives. Multiple applications exist within each vertical, fostering open competition [106, 46][Paper 1: 61]. Applications and the underlying smart contracts are permissionless, allowing applications to openly access other applications and create arbitrarily composable financial services. Composability and competition represent core values in blockchain-based finance [65][Paper 2: 121]. Aggregators serve to connect services from several applications, either vertically to provide the best rates and execution within a single sector or horizontally by combining multiple sectors to offer a one-stop solution [106, 8]. User-facing web apps or wallets can access all layers of the DeFi stack via remote procedure calls, abstracting away complexity and creating an experience akin to

¹<https://defillama.com>, accessed 16th May 2023

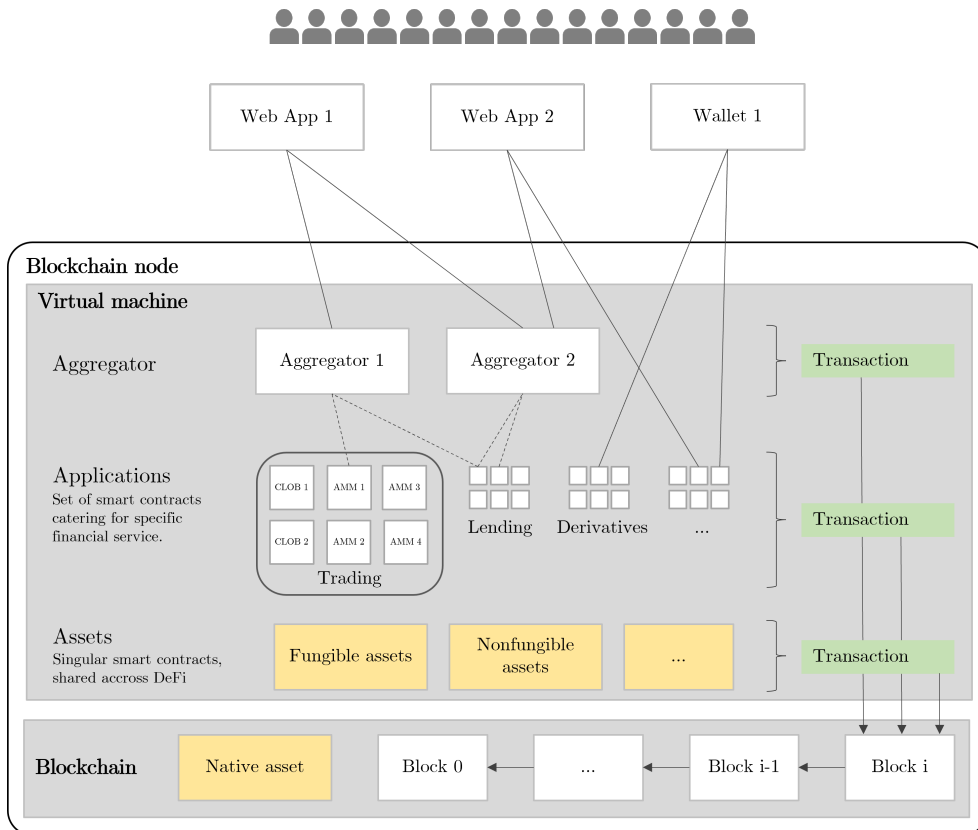


Figure 2.2: DeFi can be conceptualized in layers. The blockchain is the lowest layer and serves as a settlement layer for state-changing transactions powered by its native protocol asset. The upper layers are empowered by a VM on top of a blockchain. The DeFi application layer is formed by applications that are implemented as smart contracts catering financial services. These services can be organized in verticals, such as trading, lending, insurance, and derivatives. The figure is informed by [Paper 1: 61].

traditional banking services [Paper 1: 61].

Notably, each upper layer's security is contingent on the layer below it - the layers are hierarchical [106]. If a lower layer is compromised, the upper layer becomes unsecured. Decentralizing an application is pointless if the underlying blockchain is centralized. Thus, a decentralized blockchain is a prerequisite for a decentralized application. Further, utilizing a decentralized base layer does not automatically guarantee a decentralized application [106]. Exposed to all layers of the DeFi stack, application developers must take deliberate decisions on a range of phenomena typical to blockchain and DeFi. This requires a nuanced understanding of blockchain technology's technical properties as well as the unique characteristics and values of DeFi.

The abstract principle of *decentralization*, implying that DeFi applications are not controlled by a single entity, is typically presented as a key objective of DeFi [138, 106, 120]. This idea manifests through the underlying blockchain technology, where transactions are technically decentralized in a distributed network and are organizationally decentralized owing to a lack of a central authority [55, 133]. Applications add another layer of decentralization. For instance, applications may grant special privileges for certain addresses, enabling these addresses to pause smart contracts or to veto or alter configurations. Special privileges pose both benefits and risks [5, 104]. They can be beneficial for swiftly addressing unexpected issues. Conversely, they can be dangerous in case of power abuse or compromise [139]. To mitigate these risks, developers have the choice to either introduce special privileges in a decentralized application or design a smart contract to be truly independent. Further, the potential of abuse is often reduced by the introduction of multi-signature schemas, which necessitate *m-out-of-n* keys for a valid signature. This means that access to these privileges is dispersed across n accounts, with the signatures of m accounts required to authorize a change [5, 106]. Centralization and decentralization sit on a continuum. While decentralized applications are initially more centralized, allowing developers to swiftly iterate their design and implementation, over time, they seek to accomplish greater decentralization so as to enhance the protocol's resilience [120]. Ultimately, the pursuit of decentralization in DeFi seeks to create a more open, democratic, and resilient financial system [51, 23].

The programmability of smart contracts enables the issuance of new assets beyond the native protocol asset [5, 104]. Technically, these assets are represented by digital tokens, which are created and maintained by a smart contract. Tokens are very versatile and are capable of representing various assets such as currencies, shares, votes, and loyalty points [Paper 5: 49]. The development of token standards has emerged to foster interoperability among DeFi applications and user-facing applications, acting as interfaces to standardize the semantics of central abilities such as creating, tracking, and transferring tokens [Paper 5: 49]. Token can be broadly categorized into fungible and nonfungible token (NFT) [106]. Fungible tokens, which are interchange-

able assets, are built on Ethereum typically using the Ethereum request for comment (ERC)-20 standard. In contrast, NFTs represent unique digital objects such as photos, audio, or video [86, 102]. NFTs ensure the indisputable identification and precise tracking of the current and past ownership of digital objects. NFTs are built on Ethereum typically using the ERC-721 standard.

Besides their role in representing assets, tokens are central to the distribution of voting power among application stakeholders in token-weighted governance models [120, 115]. Governance in DeFi involves decision-making processes related to protocol changes or upgrades, usually carried out by stakeholders voting on proposals. Stakeholders express their opinions via simple majority voting schemes, with the weight of their vote directly correlated to the balance of their governance tokens. Token-weighted governance model grants users of DeFi applications a direct influence, fostering a sense of community ownership [140, 15]. The principle inherent in such a model is that, owing to their substantial investment, larger stakeholders are strongly incentivized to contribute high-quality governance. Another special token types, *stablecoins*, are among the most highly capitalized forms of crypto assets. These digital assets are typically pegged to traditional currencies such as the USD, providing a bridge between the decentralized and traditional finance systems [85]. Because of their greater stability, they are often used as a means of exchange or store of value in DeFi applications. Stablecoins facilitate seamless global transactions and are a core integration in many DeFi applications [67, 66].

A defining characteristic of DeFi applications is the high degree of interoperability [65, 47, 8][Paper 2: 121]. Owing to the unified state and deterministic execution environment shared by smart contract blockchains, DeFi applications are highly composable. DeFi essentially functions as a single, coherent market [109]. Smart contracts can invoke one another, creating novel and arbitrary complex services [65]. Through the promotion of interfaces via open-source development, applications can connect to anything created before, and provide ever-expanding services to a multi-component financial system [47]. From a technical perspective, [65] defined DeFi composability as a program that levers at least one account belonging to another program within a single transaction to create a novel financial service. For consumers, an interoperable and competitive DeFi ecosystem is desirable since it creates return opportunities and increasingly exotic financial instruments [Paper 1: 61]. However, this advanced integration also creates dependencies among both assets and applications, rendering the system susceptible to economic and technical risks [Paper 2: 121]. Like their counterparts in traditional finance, DeFi applications do not exist in isolation. Financial integration in DeFi draws numerous ideas from systemic risk research in traditional finance, a discipline that seeks to answer fundamental questions about how integration contributes to system fragility and how shocks propagate through financial systems [20, 41, 11]. Financial contagion occurs when a shock that affects one application spreads

to the rest of the system, with a single vulnerability potentially propagating across the network and affecting stakeholders throughout the ecosystem [47]. Tightly integrated financial applications can contribute to a fragile financial system in the event of a shock [20, 41, 11]. This advanced integration can occur on several layers of the DeFi stack, including the settlement layer between two blockchains, among assets [Paper 2: 121] on the protocol layer, or among applications [65] on the application layer. Examining systemic risk in DeFi is particularly important as blockchain technology inherently encourages integration, and neglecting this aspect could lead to a limited understanding of financial markets in extreme scenarios [109, 107, 65]. Notably, a dangerous scenario is the potential vulnerabilities that can arise at the interfaces through protocol interaction [110], given that protocols are rarely designed with interdependency in mind.

In recent years DeFi, has surged in popularity, offering a variety of financial services on blockchain technology in an environment that is simultaneously open, competitive, and adversarial [23, 25]. The layered conceptualization of DeFi demonstrates the intricate dependency of each layer’s security on the one below it, forming a complex ecosystem. It is not easy to design applications in this ecosystem; it necessitates a multifaceted understanding that blends technical expertise, economic design, and a social perspective, all attuned to the unique principles and values intrinsic to DeFi [133, 132]. In sum, as DeFi continues to evolve, the need for and the importance of this nuanced understanding will only increase.

2.3 Properties of Blockchain Technology and Characteristics of DeFi

This section concludes the foundation and summarizes the technical conditions that inform financial services on a public blockchain. It distills the properties of blockchain technology as well as the characteristics and values that distinguish DeFi. Numerous blockchains are available today; while distinct in their respective philosophical approaches, they share several fundamental technical properties [132, 70, 138, 133]. Understanding these shared properties, along with DeFi’s characteristics and values, lays the foundation for the forthcoming applicability framework. Table 2.1 summarizes the most relevant properties, with each property accompanied by a brief explanation and additional references for further study. As the literature on blockchain’s technical properties is extensive, these properties are described briefly and further literature is pointed to. Table 2.2 is more comprehensive and depicts the characteristics and values of DeFi as a financial application layer.

Owing to the described layered model, blockchain technology’s properties and DeFi’s characteristics are often in a hierarchical relationship [106]. DeFi’s characteristics are contingent on the blockchain layer’s soundness. If

Table 2.1 Properties of blockchain technology.

Property	Description	Literature
Decentralized	Technically, a blockchain is a distributed database and is organizationally not controlled by a central entity.	[55, 133]
Trust-minimized	A blockchain operates on a complex, formalized protocol so as to increase confidence in the computational environment. Each transaction is verified and stored by all participants. Updates that change the protocol's rules require social consensus.	[27, 52]
Immutable	Once confirmed in a block, transactions can neither be deleted nor altered. A blockchain provides a permanent record.	[132, 70, 133]
Transparent and data-rich	Transactions on a blockchain are public. Past and present transactions are accessible. Each transaction is timestamped.	[70, 133]
Pseudonymous	Transactions are associated with a specific address, which cannot be linked to a specific individual.	[136, 35, 50][Paper 3: 122]
Deterministic and Atomic	Computing transactions always return the same result. Transactions either succeed or fail completely, without an intermediate state.	[70, 4, 99]

the blockchain layer's integrity is compromised, this can put reliant applications' safety at risk. Thus, the selection of robust and dependable blockchain technology is a fundamental step toward realizing the DeFi ecosystem's potentials. Further, these categories are interdependent, not isolated. For instance, the decentralization property plays a pivotal role in increasing confidence in the computational system, and openly accessible smart contracts are critical for DeFi applications' composability.

Decentralized Technically, a blockchain is a distributed database and is organizationally not controlled by a central entity [55, 133]. Blockchain technology operates in a distributed way on a global network of computers. The individual participants are economically incentivized to operate the system. This setup results in a democratic system that is more resistant to censorship and more resilient to failure [70, 132]. However, decentralized systems can prove costly owing to challenges relating to scalability, efficiency, and governance, since coordination among a vast, decentralized network of participants can be complex and resource-intensive [70, 133, 132].

Table 2.2 Characteristics and values of DeFi.

Characteristic	Description	Literature
Endogenous	DeFi applications can only access the information present on the blockchain.	[32, 2]
Non-custodial	Users maintain control over their assets at all times.	[124, 5]
Programmable	DeFi is based on smart contracts, that can implement any logic and create novel financial services.	[70, 5]
Accessible and competitive	DeFi is accessible to any user. The applications are also permissionless such that other applications can integrate openly.	[23][Paper 1: 61]
Open-source	DeFi is strongly connected with the broader principle of open-source software development.	[133, 1]
Composable	DeFi applications have a high degree of interoperability, resulting in an integrated financial market.	[121, 65, 47]
Incentive-based	An application must strike a balance between incentivizing usage and preserving stability. DeFi is highly susceptible to market forces.	[25, 99, 47]

Trust-minimized Blockchain relies on cryptographic protocols and game theory to foster confidence in the computational environment’s operations [27, 52]. Each transaction on the blockchain is verified and agreed on, following a formalized protocol. All participants store the data, creating sufficient redundancy and thereby reinforcing the system’s integrity. Notably, any updates that alter the protocol rules require social consensus [140, 79], further enhancing the system’s credibility. While it does not fully eliminate human intermediaries, it reduces the reliance on them and ensures that no single entity holds too much power [44]. It also necessitates confidence in the technology, including the correct operation and governance of the protocol [27].

Immutable The property of immutability is central to blockchain technology. Blocks are connected through a cryptographic hash function, which references the preceding block. Altering a confirmed transaction would require changes to all subsequent blocks, a task that is prohibitively expensive, ensuring high data integrity and resistance to tampering [132, 70, 133]. However, this immutability is not immediate, as recent blocks can still undergo changes. The cost to alter increases with the transaction’s age, thus in practice, a transaction that is a few blocks old is deemed practically immutable [137, 138]. Another challenge is that this immutability applies to all transac-

tions, even those made in error, necessitating careful design and usage owing to the irreversibility of potential errors.

Transparent and Data-rich Blockchain technology is characterized by transparency and rich data [70, 133]. A blockchain publicly records transactions, providing access to both past and current transactions. Each transaction is also timestamped, creating an immutable audit trail that aids compliance, provides valuable data for market research, and allows real-time monitoring of applications and financial systems. However, the double-edged sword of this transparency is the potential compromising of privacy, since transaction data visible to all may unintentionally disclose sensitive personal information [70].

Pseudonymous The unrestricted access to the blockchain implies that any participant's balances are public. Privacy is maintained through the use of public key cryptography, providing participants with a pseudonymous address that, barring public disclosure, cannot be directly tied to a real-world identity [5][Paper 1: 61]. Thus, blockchains are pseudonymous as they allow network interaction via addresses rather than real identities. While this keeps individual transactions anonymous, it poses a challenge for regulation enforcement, as true identity can be concealed. This presents a fascinating paradox - while pseudonymous identities protect privacy, they also make illicit activities easier to facilitate [Paper 3: 122]. One limitation of this pseudonymity is its non-absoluteness since sophisticated analysis techniques can sometimes de-anonymize users by examining transaction patterns or linking addresses to real-world identities through off-chain data [64]. Generally, pseudonymity and privacy are key research areas in blockchain technology [136, 35, 50].

Deterministic and Atomic A blockchain is a deterministic system because a blockchain's state is a function of the sequence of the validated transactions that have taken place [5, 4]. Given the same sequence of inputs, a blockchain will always arrive at the same result. This principle also holds true even if the execution environment processes more complex, general-purpose transactions [5]. This deterministic property is vital for achieving decentralized consensus since each node can independently process the same transactions in the same sequence and reach the same conclusion. Nonetheless, this necessitates agreement from all the participants on the set and the order of transactions for each block, a complex process that could slow transaction speed and could induce scalability issues [70, 133]. Further, transactions are atomic, i.e. they either succeed or fail completely, with no partial states [99]. If any condition along the transaction path is not met, execution halts and the transaction is abandoned. While this ensures consistency, it also potentially leads to inefficiencies and a high demand for computational resources.

Endogenous DeFi applications can only access the information present in current or past blocks. When applications necessitate external information - such as stock prices or currency exchange rates - DeFi depends on external input [32, 2]. However, the trustworthiness of external data does not match that of the blockchain, posing a risk of introducing malicious data into the system. While a blockchain can verify the integrity of its data, it cannot verify the integrity of external information. Applications that are excessively reliant on external data could ultimately be controlled by the entity that provides the data. While solutions such as oracles exist to enhance the trust in external information [32, 2], these solutions can introduce other attack vectors and minimal reliance on external information bolsters the resilience of the DeFi ecosystem [139, 124]. Further, DeFi applications cannot self-execute and are ultimately executed as a result of a transaction invoked by an EoA [5].

Non-custodial The non-custodial attribute is another value of DeFi. Accounts are controlled by their corresponding private keys [124, 5]. Users maintain control over their assets at all times, irrespective of the DeFi application in use. This stands in contrast to traditional financial systems where central entities, such as banks, can exercise custody over assets. Non-custodial operations reduce the risk and dependency associated with intermediaries and boost privacy since account creation and usage do not necessitate information-sharing [124, 5]. However, this control comes with increased responsibility. Operating a private key requires knowledge and care, as there is no recovery mechanism when losing a private key. DeFi's non-custodial principle shifts the balance of control towards the user at the expense of increased responsibility.

Programmable DeFi is often characterized as *programmable*, since it uses smart contracts to implement any conditional logic [5]. Smart contracts enable the automation of financial transactions and programmability to create novel or customize existing applications. This led to a wide range of financial services, that could be challenging or unfeasible to implement in traditional finance [106, 51]. Although automation can potentially eliminate intermediaries and increase efficiency [70, 106], applications need to be meticulously programmed. Owing to the blockchain's openness, any programming vulnerabilities can be exploited by adversarial actors [139, 124]. Further, experimentation with smart contracts facilitates increasingly exotic financial instruments, adding layers of risk to understanding and managing these instruments. Experimental applications are particularly prone to vulnerabilities since both the smart contract and the economic interaction with the DeFi system are unprecedented.

Accessible and Competitive Built on a public blockchain, DeFi is permissionless; anyone with Internet access can participate and use its services [133].

This is a key differentiator of permissioned blockchains, which can impose access restrictions based on geography or nationality. On a more technical level, smart contracts are also permissionless, allowing applications to openly integrate with other applications and thereby creating composable financial services [Papers 1 and 2: 61, 121]. Open access potentially fosters a more equitable system, in which participants have equal access and can compare the offered services, directly fostering competition among the applications [23].

Open-source DeFi has a strong relationship with the broader principle of open-source software development [133, 1]. In technical terms, smart contracts are stored as bytecode on the blockchain. Without voluntary publication, there is a limited possibility of deducing a program’s logic from this bytecode [5]. Only upon voluntary publication of the source code, it becomes possible to match it with the deployed version. Many applications prefer to open-source their source code, providing public access for inspection, copying, and modification, a practice that ultimately creates trust in a financial service [5, 106]. With technical capabilities, anyone can verify an application’s integrity. Open-sourcing benefits the DeFi ecosystem in three primary ways: it allows for the technical verification of transactions processing and vulnerability detection in applications, and it standardizes common functionalities such as access, ownership, and voting, boosting interoperability and code security. Lastly, it encourages competitive yet collaborative development, as applications can be modified and innovated upon. While open-sourcing code is voluntary, many DeFi applications adopt this practice, given the fundamental trust it builds in a financial service [106, 5]. However, this transparency also exposes the code to potential attackers, who can exploit vulnerabilities, leading to several attacks on DeFi applications owing to public vulnerabilities in their smart contracts.

Composable A defining characteristic of blockchain-based finance is its composability. Applications have a high degree of interoperability owing to a unified deterministic state and execution environment, permissionless access, and the promotion of open-source interfaces [65, 47][Paper 2: 121]. DeFi applications can seamlessly integrate with and build on other applications, creating novel, ever-expanding financial services. Applications can utilize existing applications without needing to recreate basic functionalities. This leads to a multi-component financial system [47], which functions as a single, coherent market [109]. Users benefit from this interoperability since they can seamlessly move their assets between different applications, thereby finding customized financial services that cater to their preferences [Paper 2: 121]. Further, an interoperable and competitive DeFi ecosystem optimizes the distribution of assets and liquidity under strong market forces. However, this advanced integration creates dependencies among both assets and applications, rendering

DeFi susceptible to economic and technical systemic risks [Paper 2: 121, 124]. If one application fails, the vulnerability could potentially cascade and affect other applications. It further results in an increasingly complex financial system, making it harder for all the participants to understand the full implications of interacting with certain applications. While composability is a key driver of innovation and competition in DeFi, it also introduces layers of risks that must be carefully managed.

Incentive-based Incentive mechanisms have a key role in DeFi, given the competitive and untrusted nature of this open environment [25, 99, 47]. Applications design should seek to strike a balance between attracting users and preserving an application’s overall security and efficiency. Notably, incentives must be constructed to uphold an application’s equilibrium, as the delicate balance between supply and demand, a state that fosters system stability and ensures the effective functioning of the market. This equilibrium is often achieved via the deployment of sophisticated economic mechanisms [51]. Poorly designed economic incentives can have catastrophic consequences - from exploitation and application irrelevance to a potential collapse of the entire DeFi system [18]. Arbitrage opportunities arise in DeFi when price discrepancies occur between different applications [99, 25][Paper 6: 53]. Arbitrageurs can exploit these price differences by simultaneously buying and selling the same asset on different applications [Paper 6: 53], profiting from the spread. In doing so, they also contribute to the DeFi ecosystem’s overall efficiency by helping to maintain price consistency across markets and reducing price discrepancies [99, 25][Paper 6: 53].

In sum, public blockchain technology with smart contract capabilities, offers a trust-minimized, immutable, and highly composable system. The blockchain layer largely dictates the technical properties of the consensus protocol, the data structure, and the execution environment for smart contracts. Ultimately, these technical properties establish confidence among all participants in the computational operations. The application layer extends these properties: DeFi is shaped by core characteristics such as programmability, composability, and accessibility. Further, DeFi adheres to core values such as open-source development and a strong reliance on economic incentives. DeFi employs economic incentives to maintain a balanced ecosystem of financial services.

Financial services on blockchain technology are presented with a unique environment that is simultaneously open, competitive, and adversarial. DeFi offers continuous operations, with transactions executable at any time and accessible via an Internet connection. The principles of open access and composability allow users to smoothly transfer their assets between applications, fostering a competitive environment that drives innovation as applications

strive to self-differentiate and deliver value to users. However, this pseudonymous and open-source nature can attract adversarial actors, emphasizing the need for sound technical and economic design to mitigate vulnerabilities in applications.

Blockchain technology presents compelling technical conditions that promote a revaluation of financial service infrastructure. Yet, for applications the technology can be simultaneously enabling and limiting. While it permits the creation of innovative financial services, it also imposes certain restrictions on its capabilities. Thus, designing applications for blockchain technology is a complex task, owing to the novelty of the technology. This task requires a multifaceted understanding that blends technical expertise, economic design, and human interactions, all tuned to the blockchain's unique properties and DeFi's characteristics and values. Further, not all financial services may be suitable to implement with blockchain technology raising the question of the extent to which blockchain technology can be utilized for financial services.

Research Approach of the Thesis

This thesis is a culmination of three years of research, consisting of an essay and six individual academic publications. The publications were the focus of my Ph.D. program. Starting with an overview in Table 5 they are attached in their published form in the second part of the thesis. Five of these publications have undergone peer review, while one is still in the submission phase. All papers, although distinct in their focus and method, contribute to a nuanced understanding of the potentials and limitations inherent in blockchain technology facilitating financial services. The essay consolidates these findings into distinct items, forming the basis of an assessment framework. The framework proposes 24 guiding principles that determine when the use of blockchain technology is advantageous for a financial service.

The essay's research question is motivated by the significant potentials of DeFi as well as remaining challenges to develop DeFi applications due to the distinct environment [23, 25], nascent technology, and inherent trade-offs [133, 132]. Thus, this essay explores the extent to which blockchain technology is suitable for the implementation of financial services. The essay starts with a foundation on blockchain technology and DeFi. To facilitate a fundamental understanding of the capabilities, the essay continues with technical conditions that inform financial services on blockchain technology. The properties of blockchain technology, as well as the characteristics and values of DeFi, are drawn from the first publication [Paper 1: 61] and the underlying literature therein [23, 70, 130, 106].

The essay culminates in a framework that assesses when the use of blockchain technology is advantageous for a financial service. First, the framework is positioned in the existing academic literature. Then, the framework is presented in 4.1, it consists of 24 evaluation items, that are synthesized from the other five publications [Papers 2 to 6: 62, 49, 122, 121, 53], again supplemented with academic literature. If a publication informs an evaluation item explicitly, it is cited in the last column in Table 4.1. Afterward, the framework is briefly evaluated by applying the financial service implemented in [Paper 4: 62]. The evaluation seeks to exemplify how to apply the assessment framework to a financial service. Then, relevant limitations of the assessment framework and general chances and limitations of DeFi are discussed. Figure 3.1 depicts

the described research approach.

Categorized by their shared domain DeFi, the publications apply two distinct research methodologies, reflecting the trajectory of my Ph.D. journey from a technical to an analytical perspective. The earlier works [Papers 4 and 5: 62, 49] explore the potentials of DeFi through a Design Science Research (DSR) methodology, developing prototypes to address a specific issue at hand. The developed prototypes serve as proofs of concept and enhance the technical understanding of the problem [56]. Later publications [Papers 2, 3, 6: 122, 121, 53] adopt network analysis and analyze the data-rich DeFi environment to understand the de facto uses of DeFi. The empirical publications seek to broaden the perspective from purely technical aspects to including human interaction and economic incentive design [119, 86]. One publication notably bridges this journey [Paper 4: 62], developing a prototype and using network analysis for its evaluation.

3.1 Publications

The publications included reflect the attempt to progressively increase my contribution to each paper during the Ph.D. process. The invaluable feedback and guidance provided by my supervisors supported this journey. In earlier works [Papers 1, 4, 5: 61, 62, 49] I learned significantly from the collaboration with an experienced research team, which already had developed ideas, concepts, and methodologies. The later publications are marked by my personal ideation, initiation, and subsequent corresponding authorship [Papers 2, 3, 6: 122, 121, 53]. The later publications include two instances where seeking complementary co-authors became necessary [Papers 2 and 6: 122, 53]. The experience of collaborative writing was very valuable, since I sincerely appreciate the diverse perspectives and backgrounds of my research colleagues.

Three work-in-progress papers, that happened at a very early [60, 120] or late stage [63] in the Ph.D. program, have been intentionally excluded from this thesis, owing to their preliminary state and absence of peer review.

Given the exploratory and non-sequential nature of academic research, the six included publications are organized by their final publication date rather than their appearance during my Ph.D. program. A brief summary of each publication, along with its contribution to this consolidating essay and to my Ph.D. process is provided in the following paragraph:

1 An Introduction to Decentralized Finance (DeFi) This publication is an introduction to the concept of DeFi. It theoretically positions DeFi within blockchain’s technological context. It provides a taxonomical overview of DeFi’s applications and participants and identifies key risks. The paper contributed an early conceptual introduction to DeFi and raised relevant research areas. This was the first publication of my Ph.D. process and was valuable in

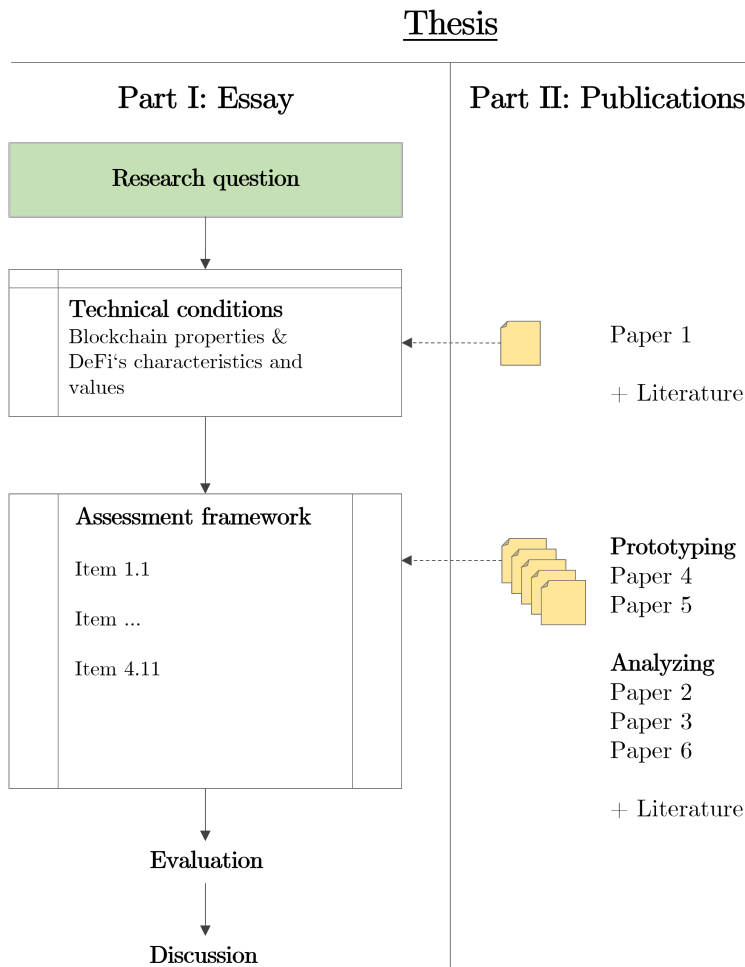


Figure 3.1: The figure depicts this essay’s research approach. It starts with motivating and posing the research question. For a fundamental understanding of the technical conditions that guide financial services on blockchain technology, the essay continues with the technical properties as well as the characteristics and values of DeFi. Both are informed by the introductory publication [Paper 1: 61] and additional academic literature. The essay derives a framework that assesses when the use of blockchain technology is advantageous for a financial service. The framework generalizes the insights from the other five publications [Papers 2 to 6: 62, 49, 122, 121, 53], again supplemented with literature. The framework is briefly evaluated and discussed, providing perspectives with relevant criticism and generalizing the challenges and limitations of DeFi. The academic publications were the focus of my Ph.D. program and are attached in the second part of the thesis. They can be categorized by their respective research methodology, either prototyping or analyzing DeFi applications.

synthesizing and conceptualizing the existing knowledge at the beginning. Led by my research colleague, I was involved in investigation, visualization, writing and editing. The paper is included in this essay, particularly for the foundation chapter introducing the layered concept of DeFi (Figure 2.2) as well as the characteristics and value of DeFi (Table 2.2). For the assessment framework, the paper contributed more generic evaluation items, such as identity and privacy, that are also part of other blockchain-focused assessment frameworks [42, 108]. Notably, the paper raises the question of composability and systemic risks in DeFi, one of the key research areas during my Ph.D. tenure.

2 Measuring Asset Composability as a Proxy for DeFi Integration

This publication explores the composability of DeFi applications. Enabled by the high interoperability of blockchain technology, DeFi applications are integrated into one another to offer attractive and novel financial services. The publication follows an asset-centered perspective, measuring how often assets are re-used in DeFi, a process akin to rehypothecation in traditional finance. The publication proposes a domain-specific algorithm to quantify the extent of these integrations. Empirical data is used to evaluate the algorithm. The results indicate a trend toward higher integration levels. Further, the data is used to hypothesize properties of financial assets that lead to higher integration levels. The publication contributes to the research on blockchain technology's financial integration and the broader research area of systemic risks. This publication was my first as the corresponding author and important in learning how to form an idea and translate it into a concept. Further, it was also the first in applying an analytical research approach, tapping into network analysis methodologies and the data-rich environment of DeFi. The publication is included in this essay, owing to the defining characteristic of DeFi's composability and the inherent risks of financial integration. The paper enhances the understanding of these and adjacent characteristics such as programmability and the data structures of blockchain technology. These characteristics are key to lever the potential of DeFi when building financial applications on blockchain technology. Further, quantifying the integration of DeFi is critical for monitoring systemic activities. Thus, the paper is also representative of broadening the scope from an application perspective to a systemic perspective. Systemic risks in DeFi remain a fascinating opportunity for future research.

3 NFT Washtrading - Quantifying Suspicious Behavior in NFT Markets

With a sharp increase in trading volume on NFT marketplaces, the open system may also be used to facilitate illicit behavior. This publication seeks to quantify the extent of such behavior and proposes two domain-specific algorithms that indicate illicit behavior on NFT marketplaces. The algorithms are evaluated against a large empirical dataset of public blockchain data. The

data is used to build highly granular transaction graphs of each NFT, since NFT's are unique digital tokens the current and past ownership is retrievable. The paper conceptualizes patterns of market abuse and their likelihood. The measured illicit behavior was lower than what industry observers estimated. However, the publication was published prior to the emergence of token incentives, i.e. mechanisms to promote trading activity. These mechanisms increase illicit behavior significantly. The paper contributes to research on the market abuse of public blockchain technology. The publication was my second as the corresponding author and valuable in increasing my methodological experience. The dataset was large and the algorithm computationally expensive, such that new ways, for instance using computer clusters, were necessary to extract, transform, store, and analyze blockchain data. The publication is included in this essay, contributing valuable insights about application misuse and thus resilient application design. Further, the publication provides in-depth empirical insights into the impact of privacy, identity, and pseudonymity on DeFi.

4 Blockchain-based Financial Infrastructure for Emerging Economies

Motivated by a large number of people in emerging economies without access to basic financial services [28, 7] and the high cost of banking services [51, 21], this publication explores the use of blockchain technology to create a prototype capable of performing three fundamental financial services: processing financial transactions, maintaining a savings account, and distributing targeted stimulus payments. First, the limitations of existing financial infrastructure for emerging economies are analyzed. The blockchain artifact focuses on targeted stimulus payments, a financial service characterized by a high demand for transaction throughput. Further, the artifact includes an asset reserve integrated into an external DeFi application, with governance mechanisms to govern these funds. Thus, the artifact explores the dilemma of requiring high security guarantees for the assets reserve while simultaneously high performance for the micropayments. Ultimately, this results in a prototype that uses and bridges two distinct blockchains: one for the asset reserve and one for performance. The prototype is deployed on the Ethereum blockchain and after a nine-month pilot phase network the prototype is evaluated through network analysis. It provides insights into the de facto uses of the prototype. During the pilot phase incentive campaigns are launched aiming to attract more users; the success of these campaigns is evaluated with the same dataset. The publication seeks to contribute theoretical and practical insights to the IS discourse on the transformative capacity of blockchain technology. This publication is an early work and is particularly representative of the Ph.D. process. First prototyping a financial service adopting DSR and subsequently monitoring the application during a nine-month pilot phase, the publication bridges both key research methodologies. Led by my research colleague, I was involved in the

concept, methodology, and writing. Further, I had the opportunity to lead the analytical part of the publication. Lastly, I presented the publication at the ECIS 2022, where the paper was a candidate for the best paper award. Owing to the limited technical capabilities of blockchain technology, the publication was early in exploring multi-chain approaches. Applying both a technical and analytical perspective, the publication contributes fascinating insights to the evaluation framework with respect to blockchain's throughput.

5 Kickstarting Blockchain: Designing Blockchain-based Tokens for Equity Crowdfunding This publication explores the boundaries of blockchain token engineering by designing, developing, and evaluating an equity token rich in features. The resulting artifact implements features common for equity such as shareholder voting, dividend payments, and documentation. The prototype relies on and contributes to several blockchain token standards. The publication culminates in seven design principles contributing to the design theory of sophisticated blockchain-based tokens. The prototype potentially implements too many features *on-chain*. Thus, the publication is well-suited for the thesis since it is representative of a key question for financial applications on blockchain technology: which features truly benefit from blockchain technology? Further, the publication enriches the thesis with insights on open-source, interoperability, tokens, standardization, and access management. The publication is an early work during my Ph.D. program and presents a major overhaul of a Master's thesis, where I was particularly involved with the concept, methodology, evaluation, and technical prototype. The publication contributed strongly to my learning of the academic publication process; the publication process culminating in acceptance at the Electronic Commerce Research Journal took 2.5 years.

6 Fundamentals of Perpetual Futures Perpetual futures are financial derivatives that never expire. Offering investors levered exposure, they are the most traded financial instrument for crypto assets by trading volume¹. Derivatives derive their price from an underlying asset. Nonetheless, for perpetual futures, it is not guaranteed that the price of the derivative converges with the underlying asset. This creates arbitrage opportunities. The publication derives theoretical non-arbitrage boundaries for perpetual futures. The boundaries are evaluated empirically by conducting a time series analysis; the publication documents considerably larger deviations than those in traditional financial markets. These deviations diminish over time, indicating a maturing crypto asset market. The paper was the last during my Ph.D. program and was initiated during my research abroad at Reichman University, Tel Aviv. I initiated the project and had significant degrees of freedom with the research idea, concept, and methodology. The paper also required a complementary

¹www.coingecko.com, accessed 15th May 2023

co-author to lead the theoretical derivation of the no-arbitrage condition. The paper is currently in the publication process. While the publication is distinct in methodology and domain, it is valuable for the essay, since it comprehensively explores arbitrage and game-theoretical mechanisms for domain-specific financial instruments. Perpetual futures are among the most sophisticated and popular financial instruments for crypto assets. Further, the publication is representative of the fascinating research collaboration abroad, where I had the opportunity to learn the different perspectives populated at an external university as well as the research area of finance.

3.2 Methodologies

Blockchain technology is a combination of distributed systems, cryptography, and economic incentives [38, 133, 70]. The computer science discipline provides the technical foundation, such as distributed systems and cryptography, enhancing blockchain's functionality, security, and scalability [133, 70]. Economics informs the design of incentive mechanisms, ensuring that all participants act in ways that foster activity while maintaining a system's stability and efficiency [13, 124, 37]. Further, the inclusion of a social perspective enhances the understanding of human interactions with DeFi applications. IS research is an interdisciplinary field, combining elements of computer science, economics, and social science [9, 105]. This research domain investigates the design, implementation, management, and use of information systems. IS seeks to formulate new theories, models, and best practices that optimize the uses of these information systems [9, 105]. By generating academic insights and practical solutions for real-world challenges, IS research holds relevance for both industry and academia. Given its interdisciplinary nature, scholars argue that IS research is well positioned to facilitate blockchain research, utilizing research methodologies that bridge technological, economic, and social disciplines [112, 90, 38].

Design science research DSR is a popular approach of highly applicable research within IS [45, 57, 81]. DSR primarily focuses on investigating a specific problem in a cyclical build-and-evaluate process culminating in a purposeful design artifact [56]. A defining characteristic of the research approach is design objectives that state requirements for a final prototype. After multiple iterations of design and development, the final prototype is evaluated against the initial objectives [95]. Figure 3.2 schematically illustrates the research approach conducted in [Papers 4 and 5: 49, 62]. This research approach has been widely adopted in blockchain research [102, 92, 104] and responds to the call for design-driven, interdisciplinary research in this field [75, 104, 103]. Building an instantiation in a domain when confronted with new technology was recommended by [57]. DSR's research output is a proof of concept and




Figure 3.2: A schematic DSR approach as conducted and presented in [Papers 4 and 5: 49, 62]. Typical for DSR are the design objectives that document requirements for a final prototype. After multiple iterations of design and development, the final prototype is evaluated against the initial objectives. This specific iterative research approach originates in [95].

is useful in understanding the problem [56]. This output may take various forms, such as frameworks, software architectures, or prototypes. [Papers 4 and 5: 49, 62] apply the DSR methodology.

Network analysis The analysis of networks has been a pivotal methodological tool in computer science to understand complex systems. Network analysis is rooted in graph theory [34] which is concerned with the properties and structures of graphs. However, network analysis expands on this foundation, incorporating interest not only in a network’s inherent properties, but also in the behaviors and interactions of the objects it represents [127, 126]. An extensive introduction to graphs and networks is provided by [88, 125]. Network analysis is a versatile method for researchers in understanding the complex DeFi ecosystem and has been increasingly employed as a research methodology in DeFi [109, 65, 22, 118, 36, 117]. During my Ph.D. program, three publications adopted a purely analytical perspective [Papers 2, 3, 6: 121, 122, 53].

This research area can be broadly divided into two categories: those proposing analytical tools, and descriptive research into DeFi applications and the ecosystem. The former aims to provide applicable techniques to address fundamental challenges in DeFi research. For instance, one study proposes an algorithm to detect address clusters for Ethereum aiming to recognize entities and enable a more precise assessment of overall network statistics [117]. [36] proposes a technique to investigate the use of standards across the entire DeFi ecosystem and how open-source code is copied and modified. Descriptive research can focus on specific phenomena in DeFi, such as tokenized assets or governance. [22] and [109] investigate ERC20 tokens and systematically compare activities and relations between tokens. [120] analyzes the distribution of voting power granted by tokens across DeFi applications and [115] investigates whether token-based governance is generally beneficial. Other descriptive research investigates application-specific aspects. [96] analyzes liquidations on a lending application and [48] compares use and configurations across several

lending applications. [86] investigates and compares the market structures of NFT marketplaces on the Ethereum blockchain. Notably, a fascinating analytical research stream extends to a systemic perspective of the entire DeFi ecosystem. [65][Paper 2: 121] seek to gauge the system's integration. Understanding integration across DeFi applications is essential to manage the systemic risks of this nascent ecosystem.

Toward an Assessment Framework

This chapter derives the evaluation items for the assessment framework presented in Table 4.1. It is organized into a section that derives items primarily informed by the two prototyping papers [Papers 4 and 5: 49, 62] as well as a section, that derives items from the analytical publications [Papers 2, 3, 6: 122, 121, 53]. The items are supplemented with relevant academic literature. For a balanced evaluation, the framework uses a score-based approach, similar to [108]. Each item poses an active question and is marked with a (+), (0), or (-), indicating the directional impact of a *Yes* answer. If a (+) item receives a *Yes* response, it increases blockchain technology’s suitability for the financial service, while a *Yes* answer to a (-) item decreases its suitability. A (0) item’s impact is not straightforward and can lean either way. These (0) items demand further discussion.

4.1 Items from Building DeFi Applications

Published during my Ph.D. program the two papers [Papers 4 and 5: 49, 62] use DSR as an applied research methodology to explore financial applications on blockchain technology, primarily through the construction of prototypes. DSR focuses on investigating a specific problem in a cyclical build-and-evaluate process [95]. The process culminates in a purposeful design artifact, serving as both a proof of concept and useful in understanding the problem [56]. For the assessment framework, prototyping presents a valuable research method. The technical insights derived from these prototypes enrich the framework through a technical lens, further supplemented by relevant literature from the IS and computer science disciplines. For instance, [132, 70, 137, 40, 133] compare different blockchain networks so as to organize the different technical aspects of blockchain technology. [5, 4] provides a technical manual for each of the two most popular public blockchains, Bitcoin and Ethereum. Finally, [12, 80, 111, 128] elaborate on smart contract development and empirically analyze design patterns in existing smart contracts.

Item 1.2: Does a well-functioning solution already exist? (-)

Before deciding to use blockchain technology, one should examine existing solutions and should validate that a significant problem remains unaddressed [42]. For instance, [Paper 4: 62] examines existing mobile payment solutions in emerging economies and motivates a blockchain solution with relatively high costs for the end user [21]. The question facilitates an assessment of existing solutions and their limitations, and whether the problem at hand has previously been addressed using alternative technology. The item is rather strict in that blockchain technology should be explored if an existing solution is well-functioning except for only one critical limitation. Blockchain technology's properties can be used to potentially address a wide variety of limitations in existing solutions, such as a centralization of power [70, 132], high costs [Paper 4: 62] or a lack of transparency.

Item 1.1: Is the service significant enough to justify the costs of a decentralized network? (+) Another important consideration is whether a financial service justifies the costs to participate in a decentralized network, for instance through a high user count or large transaction volume. Building a financial application on a decentralized network comes with significant development and operational costs, compared to traditional, centralized server-based applications. Owing to its design, where every state change is sent, verified, and replicated across all network participants, blockchain technology inherently requires more resources and network communication [132, 133]. Thus, public blockchains operate slower compared to centralized databases, precisely because of the requirement of synchronizing resources among multiple, non-trusted participants [93]. This question ensures that the increased costs are proportional to the service's significance. While some services may justify the costs of decentralization, such as targeted stimulus payments for multiple countries [Paper 4: 62], other services may not. For instance, a private service without a direct financial value such as a point-based loyalty program. Being aware of the costs of decentralization helps to avoid unnecessarily expensive solutions to trivial problems or overlooking more cost-effective solutions to a problem at hand.

Building a DeFi application on blockchain technology presents several technical restrictions.

Item 3.3: Does the service require high throughput? (0/-) First, the transaction processing capacity is limited [24, 77] as it depends on the number of transactions per block and the block frequency, determining how many transactions a blockchain can process per second. Transaction capacity varies greatly across networks, with Bitcoin processing around five TpS and Ethereum slightly over 25 TpS. High demand - especially in a short amount of time - financial applications can be limited by these throughput rates. [Paper 4: 62] addresses this issue by utilizing two blockchains. The asset reserve and

governance components, both with high trust requirements, are hosted on a secure blockchain. The micropayment component with performance requirements is hosted on a high-performing blockchain. Further, congestion on the blockchain can lead to increased transaction fees, affecting the affordability of using blockchain applications [32]. Notably, scalability has been a key research area [24, 77], and there are ongoing efforts to increase the transaction rate of blockchains. Gauging the throughput requirements of a financial service is a recurring question posed by several assessment frameworks [78, 69, 42].

Item 4.5: Does the service require complex computation? (-)

Second, to prevent hostile activities, infinite recursion, or excessive use, computational operations on a blockchain are limited and priced. There are two outcomes: the sophistication of financial applications is restricted and computational operations incur costs. This limitation of operations results in DeFi applications only being near-Turing complete [5]. Thus, blockchain technology is not well-suited for heavy computations or sophisticated algorithms, especially if these are needed promptly [132, 5]. The computation's sophistication must be evaluated against the costs and limitations of the blockchain network. [Paper 5: 49] introduces an equity-like token, with typical shareholder rights such as dividend payments and shareholder voting. By potentially implementing too many features, this paper explores token engineering's boundaries and highlights the necessary *on-chain-off-chain* decisions in DeFi applications. On-chain-off-chain decisions identify aspects of a financial service that benefit from a trusted and decentralized execution environment and those that can be implemented in a cheaper traditional and centralized technology.

Item 3.2: Does the service store many nontransactional data in each transaction? (0/-) Further, data storage in blockchain technology is inherently expensive and less efficient than traditional databases, with all network participants permanently storing the entire blockchain [70, 5]. Large volumes of non-transactional data can lead to rapid blockchain growth, making it increasingly resource-intensive to maintain and synchronize. Therefore, when designing a blockchain application, it is crucial to consider the nature and size of the data that each transaction will carry [132, 78, 69]. Applications can outsource the expensive operation of storing non-transactional data to decentralized storage technology [71] using *off-chain* alternatives. For instance, [Paper 5: 49] utilizes a distributed file-sharing system (IPFS¹) to store documents required for equity-like tokens. The need for data outsourcing is further exemplified by the emergence of NFTs, where the data-intensive art pieces are often stored on IPFS [86].

Item 2.4: Does the service benefit from strict control and a

¹<https://ipfs.tech/>, accessed 13th June 2023

majority-based approach for protocol updates? (+) Another consideration when building a DeFi application is that blockchain technology is still nascent and subject to significant updates. Fairly recent updates have addressed central elements such as consensus² and cryptographic signature schemes³. Since blockchains are decentralized networks, updating the existing rules of a blockchain protocol requires agreement among the majority of the participants, also known as social consensus [140, 79]. Updates are typically implemented as a *fork*: a divergent blockchain that follows the new protocol rules [133]. For a successful update, the majority of participants must switch to the new blockchain at a specific time. The difficulty of protocol updates means that most upgrades are rigorously debated, and software updates cannot be implemented by a single entity. No single entity has centralized power and can dictate a software upgrade's direction [93, 69, 42]. Otherwise, for applications, there is the risk of unfavorable or slow updates.

Item 3.4: Does the service benefit from blockchain technology's specific data properties, particularly immutability and transparency? (+/0) Data stored on a blockchain are subject to specific data properties. Transactions on the blockchain are immutable and transparent. Once processed in a block, transactions cannot be deleted or altered and are retrievable by anyone [70, 133, 132].

Financial services that use blockchain technology must account for immutable transactions. Once processed, transactions cannot be reversed, even those accidentally directed to incorrect addresses. On multiple occasions, immutability has led to accidental losses [96][Paper 1: 61] of significant financial assets. Further, the immutability also extends to the applications themselves. Once deployed, program updates, even those aimed at bug fixes or functional enhancements, are not possible. The strict immutability of blockchain technology benefits financial services that are stable in their development or that require strong auditability for compliance reasons. Financial services that require frequent data modifications or deletions may be at a disadvantage [68, 94, 78, 42]. Nonetheless, there are various design principles for sophisticated upgrade management that DeFi applications can choose to implement from [12, 128, 80]. A DeFi application typically mitigates immutability by implementing a proxy pattern, akin to a reverse proxy in traditional software design [5, 104]. A proxy pattern requires the collaboration of two smart contracts, one providing the logic (*back-end*) and another forwarding traffic (*proxy*). The proxy contract has its own unique address, and users and other applications interact solely with this proxy. During an upgrade, the old logic contract is replaced by a new one, and the proxy contract is altered to point

²<https://ethereum.org/en/roadmap/merge>, accessed 25th May 2023

³<https://hackernoon.com/bitcoin-taproot-a-multifarious-upgrade-to-the-bitcoin-ecosystem>, accessed 25th May 2023

to this newly deployed smart contract. This pattern has also been explored by [Paper 5: 49].

A blockchain is designed to provide a transparent and ordered record of all past transactions, thus offering significant advantages to a service that benefits from a transparent and data-rich environment. The technology ensures strong accountability by providing an immutable and transparent audit trail for regulatory compliance, potentially reducing disputes and the need for third party auditors. Further, the data-rich environment can support entities engaged in market research or real-time monitoring of applications and financial systems [119]. However, this transparency level may present challenges for financial services that seek less openness. It reveals sensitive information or business intelligence that would typically be kept private in a traditional setup, potentially compromising competitive advantage.

Lastly, the data are stored in copies across participants' servers, thus offering a high degree of redundancy. The redundant storage potentially increases applications' ability to withstand censorship or attacks [133].

Item 4.4: Does the service benefit from programmability and automation? (+) Building applications with blockchain technology potentially unlocks a large design space for novel financial services as well as the automation of existing ones [106, 51]. These potentials are primarily introduced through the programmability of smart contracts. These autonomous contracts can implement any conditional logic with codified instructions such as automating financial transactions or creating novel financial services [5]. Financial services can benefit from programmability, whether through increased efficiency, cost reduction, or the offering of novel services to users. Applications can use the large design space for financial innovations that would be hard or impossible to implement in traditional finance. However, applications need to be meticulously programmed because, owing to the blockchain's openness and immutability, any bugs or vulnerabilities in the program can be exploited by hostile actors [139, 124].

Item 4.8: Does the service benefit from an open-access environment? (+) Smart contracts deployed on public blockchain technology are shared across the network, and accessible to anyone for interaction [5]. Thus, applications can attract a global user base, regardless of location or identity, with only a unique address required for identification [Paper 4: 62]. The relationship between the user and an application is often a temporary, ad hoc one with users interacting for a specific purpose and then disengaging [Paper 4: 62]. Open access potentially minimizes user discrimination based on factors such as geography, social status, or nationality, enabling applications to reach a global user base. [Paper 4: 62] is exemplary of access to a global user base; during the nine-month trial period, the application processed more than 6.6 million transactions among 189,370 participants, supported by automated

incentive campaigns. However, open access also comes with disadvantages for financial services which may need to regulate access to their applications in order to comply with regulatory requirements [135]. Further, it can be hard to exclude adversarial participants. If necessary, specific access restrictions can be implemented. There are various design principles for sophisticated access management [12, 128, 80]. Access control - determining which address can perform administrative tasks - is increasingly important in the DeFi ecosystem. This control introduces potent operations, such as minting or freezing assets [104]. Initially, each smart contract is owned by the deploying address. For DeFi applications, there is an inherent expectation to minimize upgradability and access, so as to minimize technical and economic attack vectors for these applications [5]. Minimizing trust requirements in the application design follows suit with an overarching objective of blockchain technology, to reduce reliance on central entities.

Item 4.6: Does the service benefit from open-source development? (+) DeFi closely aligns with the broader principle of open-source software development, allowing applications to copy, modify, and distribute source code as they deem fit [133, 1]. While this is voluntary, many DeFi applications choose to publish their code, a practice that creates fundamental trust in the financial service. Applications without public source code have difficulties attracting users since the application's inner logic is not verifiable [5]. Open-source development can be advantageous to both applications and the entire ecosystem. Community contributions or modifying existing code can accelerate software development[1]. The collaboration on and application of standards enhances a service's security and interoperability with other financial services. For instance, [Paper 5: 49] builds on a variety of token standards to implement equity tokens, creating a prototype that is interoperable with the DeFi ecosystem. However, publishing an application's inner workings can also expose potential vulnerabilities. DeFi applications have experienced several exploits owing to public vulnerabilities in their smart contract code[139, 124][Paper 1: 61]. Financial services that can lever the positive aspects of open-source development benefit from DeFi's open-source ethos and increase blockchain technology's applicability.

Item 3.1: Is the service exposed to settlement risks? (+/0/-) Similar to prolonged settlement periods in traditional finance, the time between sending a transaction and its resolution is not immediate [70, 3]. Settlement risk describes the possibility that, in this period, one party may fail to deliver, for a number of reasons, including operational issues or insolvency. Traditionally, this risk is often mitigated by using risk-reducing third parties such as clearing houses [83, 59].

Generally, a blockchain settles transactions peer-to-peer without requiring intermediaries [55, 133]. Further, transactions are atomic - they are either

settled completely or rolled back [99]. However, the precise settlement of blockchain technology and time considerations for transactions is a very nuanced research area. Blockchain technology relies on sophisticated concepts from the computer science discipline for key issues, such as the time taken for an order to be processed (*latency*) and the point at which a transaction becomes irreversible (*finality*) [3]. While different blockchain protocols share the concepts of transactions requiring time to be processed and finalized, the specific details can differ significantly between protocols. [132, 137, 40] provide an overview over latency and finality for different blockchain protocols.

The evaluation of settlement risk gauges the tolerance for latency for a specific financial service [42]. Despite recent protocols promising attractive processing times of seconds or sub-seconds, the suitability of blockchain protocols for services that require high-speed transactions - such as real-time trading - can still be limited. Further, networks can become congested if the demand for transactions exceeds the block's limited supply [32]. While this is balanced by increasing transaction fees through an auction mechanism, these fees can react too slowly during periods of stress [30][Paper 1: 61].

Applications must further assess how critical *guaranteed* finality is for the service at hand [42]. The finality is affected by a blockchain's consensus type [3]. Bitcoin has a probabilistic finality, where transactions have a chance of being reverted, with the chance decreasing with every subsequent block. While Ethereum aims for deterministic finality after processing certain *checkpoint blocks*, these checkpoint blocks are subject to voting and complex attack scenarios⁴. In both cases, finality takes considerably longer than transaction processing, and in rare cases, finality only occurred days after the initial transaction. If a service requires extremely high speed or guaranteed finality blockchain technology can be inferior to centralized technology. For the majority of financial services, the literature suggests that blockchain technology is advantageous when a financial service faces high settlement risk within traditional finance [83, 59]. [71] indicates the potential to improve on traditional financial infrastructure in several examples, with a more efficient approach to settlement.

Item 4.1: Does the service manage digital value exchange? (+)

Blockchain technology naturally excels at managing digital value, and has certain limitations in handling physical goods [130]. Owing to the endogenous nature of blockchain technology, a blockchain can neither verify a product's physical attributes, nor whether the digital record accurately represents the physical item [2]. This disconnect is often referred to as the *oracle problem* and can lead to difficulties when dealing with physical goods [32]. Nonetheless, blockchain can offer value even in such situations when combined with

⁴<https://blog.ethereum.org/2016/05/09/on-settlement-finality>, accessed 2nd June 2023

other technologies such as tamper-proof sensors that can supply reliable data about physical goods [14, 138]. However, the requirement for supplemental systems introduces additional complexity and a potential point of failure into the solution.

The assessment framework is complemented by four items about the social dynamic and access requirements of stakeholders that are often found in other blockchain assessment frameworks. These items are fundamental to blockchain technology and can provide an early and definite indication of the applicability of blockchain technology, regardless of the application domain.

Item 2.1: Are multiple parties involved in the requirement for write access? (+) Blockchain technology establishes consensus for transactions among a set of participants. This item evaluates whether multiple entities are involved and whether each entity should have the authority to update the shared database. Building a decentralized service comes with considerable costs, both in development and operation, compared to a traditional server-based application [132]. If feasible, a centralized setup may prove to be a better solution. However, blockchain technology becomes a viable option when multiple parties are involved, each with the requirement to update the state of the ledger [58, 94, 130, 78, 69, 42].

Item 2.2: Do the parties have conflicts of trust or interest alignment? (+) This question gauges the social dynamics among the interacting parties. Blockchain technology is often referred to as a *trust-minimized* system because the parties must trust the blockchain protocol instead of one another [27, 52]. This property is very beneficial in scenarios where trust between parties is in conflict, such as in transactions among unknown parties or where interests diverge. If complete trust and aligned interests exist among the interacting parties, an alternative setup may be a better solution [58, 68, 94, 130, 42].

Item 2.3: Is a trusted third party feasible? (-) Traditional financial systems often rely on trusted third parties, such as banks and clearing houses, to manage and settle financial transactions [83, 59]. If a trusted third party can effectively mediate between multiple parties with conflicting interests, blockchain technology may be inferior. However, the requirements for a third party are high - they must offer the service truthfully, swiftly, transparently, at reduced costs, and without failure or downtime. Further, both are necessary: that the interaction parties *can* establish a trusted third party and that these parties *want* to establish a trusted third party [58, 132, 94, 130, 42].

Item 1.3: Does the service require an unambiguous regulatory setup? (-) Regulation is of great importance for financial services, though it falls outside the scope of computer science and this thesis. The essay adds reg-

ulatory aspects only for comprehensiveness, however for an extensive analysis the reader is referred to the work of [135, 91, 16, 26]. While DeFi struggles with the regulation of financial services on blockchain technology and key questions, such as the jurisdiction in case of controversy, or whether a smart contract can be liable [135], progress is being made with increasingly clear regulations emerging worldwide. Interestingly, [134, 26] and [Paper 5: 49] suggest potential for blockchain technology to facilitate regulation - the programmability of DeFi may enable 'embedded supervision' [134], where regulation could be enforced within the program code, rather than being supervised in post-transaction audits. However, if a financial service requires a definitive regulatory setup, blockchain technology is less likely to be beneficial. Conversely, if a financial service allows some regulatory flexibility, blockchain technology can be used already today [108, 71].

4.2 Items from Analyzing DeFi Applications

Blockchain technology is fundamentally data-rich, owing to its property of storing every transaction in a transparent, immutable, and chronological order. In a seminal paper, [109] coins DeFi on the Ethereum blockchain as the largest public financial dataset with a granularity of individual trades and users. This offers an improvement to common limitations within traditional financial infrastructure, where the acquisition of granular empirical data is increasingly challenging owing to regulatory constraints or intellectual property rights [109].

The ensuing section continues deriving and discussing items for the assessment framework, guided by the three analytical papers published during the course of my Ph.D. [Papers 2, 3, 6: 53, 121, 122]. These papers apply network analysis and time series studies, utilizing public blockchain data to investigate individual DeFi applications as well as the integrated DeFi system. The objective analysis of existing applications can potentially assist in both the ex ante design considerations for building an application and the ex post operation of the application. Prior to deployment, insights gained from analyzing the design and use of existing applications can inform the design and building process. After deployment, monitoring the application can provide an objective evaluation of the current use and risks, thereby guiding adjustments to configurations [119][Paper 4: 62]. Thus, the analysis of existing applications is useful in understanding the de facto uses and financial flows, and assessing whether the economic mechanisms are working as intended. This section approaches the research question from an analytical perspective, seeking to broaden the discussion from purely technical to encompass human interactions and economic design [119, 86].

Item 4.9: Does the service require centralized governance? (-)

Operating on decentralized blockchain infrastructure enables the development of decentralized applications offering financial services. This decentralization is typically presented as a core objective for building DeFi applications, seeking to reduce reliance on centralized agents and to enhance the resistance of protocols [120, 138, 10]. However, operating the application commonly requires some form of governance, which could potentially reintroduce centralized power [115, 120]. Governance becomes necessary when implementing new features, managing access permission, or altering configurations. In DeFi governance is conducted in majority-based voting through tokens [115, 120]. Notably, the decentralization of governance is not binary - the application can have varying degrees of governance decentralization [120, 10]. This item seeks to evaluate a financial service's governance requirements, as coordinating among a vast, decentralized network of participants can be intricate and resource-intensive. Typically, applications begin with a centralized design, allowing developers to iterate rapidly on design and implementation. Over time, these applications strive to achieve decentralization to bolster the protocol's resilience [120]. However, if a financial service requires continuous updates, such as active risk management, blockchain technology may be less suited. Conversely, if a service benefits from a more democratic approach to governance, blockchain technology becomes beneficial [94].

Item 4.10: Is the service suitable for game-theoretical mechanisms and strong market forces? (+) Game-theoretical mechanisms designed to incentivize beneficial and penalize adversarial behaviors are key to blockchain-based financial services. DeFi applications are often designed with game-theoretical mechanisms aimed at guiding the application toward equilibrium [51]. Applications have to strike a balance that both incentivizes usage and preserves overall stability. However, these mechanisms can be undermined by adversarial actors or unforeseen market conditions, leading to a disequilibrium that could disrupt the applications and thus a system's stability [25, 99, 47]. In a pseudonymous and open environment, actors are financially motivated to exploit systems that lack equilibrium. For instance, a service may depend on the balance between lending and borrowing rates to maintain functionality. If this balance is disrupted, it could lead to issues such as liquidity shortages or abnormal interest rates [96]. Another example includes the reliance on over-collateralization as security against loans. Loans typically require assets of equal or more value as collateral owing to users' pseudonymity. While this approach is secure, it is also capital-intensive [96, 66]. Further, user behaviors may deviate from the initially-intended design, which can be monitored by analyzing the on-chain user behaviors on DeFi applications [Paper 3: 122]. The question seeks to evaluate the reliance on and sophistication of game-theoretical mechanisms necessary for a financial service. A system that naturally gravitates toward equilibrium may be a better fit for blockchain technology than a service that is vulnerable to exploitation. Further, the fact

that applications are subject to strong market forces can be levered: incentive mechanisms can be designed to attract users and can incentivize positive activities, a phenomenon explored by [Papers 4 and 6: 62, 53]. However, economic incentives are a double-edged sword; while they can be a powerful tool to bootstrap a system, a design flaw can be catastrophic, leading to exploitation or instability [18]. Therefore, understanding whether a service tends toward a natural equilibrium can assist in navigating the complex economic dynamics of the DeFi ecosystem by anticipating potential disruptions and designing mechanisms to restore balance when necessary.

Item 4.7: Does the service benefit from composability? (+) A defining characteristic of DeFi applications is their high degree of composability, enabling the creation of a bottom-up, multi-component financial system [47]. The open accessibility property of blockchain technology extends to technical integrations by other applications, laying the technical foundation for the composability of the DeFi ecosystem [65, 47][Paper 2: 121]. DeFi applications can connect to anything created before, and create novel and arbitrary complex services. For instance, a DeFi lending application may integrate with a stablecoin application offering stable USD-denominated assets for its money market. From a user's perspective, different applications operate seamlessly together. This item evaluates the advantages of integrating with existing services on the blockchain. Composability offers advantages for applications, since it allows them to build on other services, leading to rapid development and exponential growth in functionalities [69]. Compatibility among different DeFi applications facilitates the exchange of information and assets across the ecosystem. For instance, a service could immediately tap into the liquidity of another application reducing its own liquidity requirements. However, if a financial service functions independently without interacting with other applications, blockchain technology's utility may be lessened. Further, advanced integration creates dependencies among both assets and applications, making the system susceptible to economic and technical risks [65][Paper 2: 121]. Vulnerabilities in one application can propagate and compromise other integrated applications, leading to a cascading effect through the entire ecosystem. The experimentation with smart contracts have led to increasingly exotic financial instruments, making the layers of risks involved hard to understand [65][Paper 2: 121]. Notably, the ubiquity of blockchain data provides a basis to observe dependencies across the DeFi ecosystem. [Paper 2: 121] seeks to quantify the integrations of DeFi applications from an asset perspective by analyzing shared assets and derivatives.

Item 4.2: Does the service require significant access to external information? (-) Owing to the closed nature of blockchain systems, applications can only execute operations based on information present in current or past blocks. External data - such as stock prices or currency exchange rates

- are not directly accessible, and must be fed into the blockchain [2, 32]. An application that relies heavily on external data risks being dependent on the entity that supplies the data. The external input may come from a central, malicious source, and does not have the same trust guarantees as blockchain data, posing the danger of feeding the application with malicious information. While there are technical solutions to increase trust in external information [2, 32], these solutions introduce an additional layer of potential vulnerability into a blockchain system. Minimal or careful use of external information increases the application's resilience and blockchain technology's applicability.

Item 3.5: Does the service require privacy? (+/0/-) Blockchains typically operate on a pseudonymous basis, identifying users by their public addresses rather than their real-world identities. Despite the advantages provided by pseudonymity, it still exposes a user's entire transaction history and balance owing to blockchain's transparency [5]. Users can aim for enhanced privacy by utilizing privacy-enhancing tools that can be employed to obscure transaction details while still ensuring their validity. Privacy is an extensive research area for blockchain technology [136, 35, 50]. However, even with the support of privacy-enhancing tools, privacy is not absolute. Sophisticated analysis techniques can sometimes de-anonymize users, either by tracing patterns in transaction data or associating addresses with real-world identities through off-chain data [64]. This creates a proverbial 'cat-and-mouse' game in the quest for privacy [Paper 3: 122]. Blockchain's transparency level may be inappropriate for financial services that require very strong privacy, such as wealth management [35]. In contrast, transparency can help financial authorities to prevent illicit behaviors in financial markets. Therefore, striking an optimal balance between transparency, privacy, and regulatory compliance remains one of the most challenging questions for the applicability of financial services [42, 108, 31]. [Paper 3: 122] explores this phenomenon and contributes domain-specific algorithms to detect illicit behaviors in NFT marketplaces.

Item 3.6: Which identity level is required for the service? (+/0/-) Similarly, a financial service should consider the extent of user information necessary for its operation and how significant real-world identity is for regulatory compliance [42, 31]. For instance, assessing an individual's creditworthiness in traditional finance loans requires comprehensive background information, such as proof of salary. The pseudonymity of blockchain offers only information attached to the address, making it challenging to estimate creditworthiness solely on address information [48]. [47] simplifies identity requirements for DeFi applications into three tiers: no identity, weak identity, and strong identity. A service that functions without additional identifying information (*no identity*) is better suited for blockchain technology. In contrast, services that necessitate substantial background information, or a *strong identity*, are challenging to implement with blockchain technology.

Admittedly, many financial services need robust identity verification to comply with local anti-money laundering (AML) rules. Similarly, voting systems may require identity to ensure fairness [115]. Considering the importance of identity for many applications, identity solutions for blockchain technology is an intense research area, with various external solutions available [33, 76, 101].

Item 4.3: Is the service exposed to increased counterparty risk?

(+) Counterparty risk refers to the risk that one party in a transaction could default on its contractual obligations. This is a common concern in traditional finance, especially when dealing with parties that may not be well-known or reliable. In DeFi, the counterparties are autonomous smart contracts, that execute automatically when predetermined conditions are met. Further, DeFi transactions are atomic - they are either fully executed or rolled back, without any intermediate state. Where services are increasingly exposed to counterparty risk, blockchain technology can provide a substantial advantage by ensuring that contractual obligations are automatically enforced [59, 83]. This allows for more reliable and secure transactions, even in a trustless environment.

Item 4.11: Does the service require significant conflict resolution?

(-) Although DeFi, with its deterministic smart contracts, sets predefined rules that seek to prevent conflict from arising in the first place, blockchain is less effective at resolving conflicts that involve subjective judgment or external information. For instance, while a blockchain can automatically facilitate shareholder voting for an equity token, it cannot automatically resolve disputes regarding the quality of proposals [Paper 5: 49]. A blockchain is incapable of evaluating a proposal's quality or of considering external information. For financial services that require significant conflict resolution in the form of rollbacks or corrections, especially involving subjective judgment or external information, additional conflict resolution mechanisms may be necessary [31, 108]. This ultimately introduces an additional layer of vulnerability into an application. From a technical perspective, the demand for conflict resolution can be reduced by rigorously testing applications and avoiding technical errors in an immutable environment [Paper 4: 62].

4.3 The Assessment Framework

This section consolidates the individual items into an assessment framework. Before presenting the framework, the section starts with a review of existing blockchain assessment frameworks. The final financial service-specific assessment framework proposes guiding principles that determine the degree to which blockchain technology is advantageous for a specific financial service.

To exemplify a brief evaluation, the framework is applied to the financial service of targeted stimulus payments, a blockchain prototype developed in [Paper 4: 62]. This section then discusses the framework, providing perspectives with relevant criticism, and elaborates on how the synergies between building and analyzing DeFi applications shape the assessment of DeFi applications. Lastly, this section broadens the perspective and discusses the potentials and limitations of the DeFi ecosystem.

In the IS discipline, suitable use cases for blockchain technology have long been a research topic [68, 94]. Assessment frameworks in the IS literature seek to help researchers and practitioners to determine a specific technology's suitability. Frameworks are valuable tools, since determining whether the application of blockchain is justified remains a major obstacle [94]. Thus, there have been numerous contributions to assessment frameworks that determine the need for blockchain technology [132, 58, 94, 130, 78, 42, 93, 69, 71]. The most frequent criterion in these frameworks is the assessment of the trust relationship among multiple parties [58, 94, 93, 78]. This is in line with the inherent trust-minimizing property of blockchain technology, ensuring that the trust-minimizing property genuinely provides value. Using blockchain technology can potentially improve trust issues in existing solutions. Another frequently assessed criterion is the possibility of simpler solutions. Since decentralized solutions are typically more challenging to implement, the examined frameworks refer to alternative technical solutions. This criterion can be addressed in various ways, such as checking whether a functioning system already exists [42] or if a trusted third party is a feasible option [58, 132, 94, 130, 93, 69]. Most of the reviewed frameworks propose a definitive, sequential flow diagram [58, 94, 130, 93, 78, 69]. If a question is negated, blockchain technology is entirely dismissed. A notable exception is [108], which uses a scoring approach where a series of questions are answered and the cumulative result determines the final applicability. While flow diagrams simplify the process and provide a definitive answer, a scoring approach could offer a more balanced understanding of the issue. Most of the reviewed frameworks also consider permissionless blockchain technology as an option [58, 94, 130, 93], with [69] including a broad range of other technical solutions. However, owing to its focus on public DeFi, the framework presented in this essay excludes any permissioned blockchain solutions. Composability and open access are defining characteristics of DeFi, leading to strong network effects on public blockchain technology. Bitcoin and Ethereum have attracted a large number of developers and users as well as substantial financial assets.

Despite existing assessment frameworks for blockchain technology, there remains a gap in the literature for a framework specifically tailored for financial services on blockchain technology. All frameworks focus solely on blockchain technology. [103] highlights that a blockchain implementation benefits from a balanced and individual consideration. [69] calls for the expansion of blockchain-based frameworks to include domain-specific questions, particu-

larly highlighting economic aspects. For instance, [10], proposes a framework to assess the decentralization of blockchain-based governance. This gap is remarkable, given the broad range of financial services and the rapid growth of DeFi, with billions of USDs in assets at risk. [112, 82, 43] further identify the need for research that enables applications to maximize DeFi's value while considering the inherent risk.

The framework presented in this essay consists of 24 items and generalizes the insights of the individual publications [Papers 1 to 6: 61, 62, 49, 122, 121, 53]. It also aligns with the criteria that frequently appear in the examined frameworks. The framework seeks to address the research questions of to what extent blockchain technology is suitable for the implementation of financial services.

The assessment framework is organized into four overarching categories, represented by different colors in Table 4.1. These categories are broadly influenced by the discussion that a balanced assessment requires the specification of the current state of financial infrastructure, the fundamentals of blockchain technology, technical requirements, and the utilization of DeFi values. Thus, the red *State of the art* category seeks to evaluate the current operation of the financial service and whether an alternative technology is better suited. The green *Blockchain: trust* category measures the necessity of establishing trust among multiple parties via blockchain technology and assesses whether the fundamental trust-minimizing properties of a blockchain are needed for a financial service. The blue *Blockchain: technology and performance* category identifies the technical system requirements of a specific financial service. Lastly, the yellow category *Characteristics and values of DeFi* assesses the degree to which a specific service can lever DeFi's positive potentials. Owing to their limitations on blockchain technology, the examined frameworks in the literature contribute mostly to the *Blockchain: trust* and *Blockchain: technology and performance* categories.

Table 4.1 lists each item per category and adds the relevant literature. If a Ph.D. publication informs an evaluation item explicitly, it is cited in the last column. For a balanced evaluation, this framework uses a scoring approach, similar to [108]. Each item poses an active question and is marked with a (+), (0), or (-), indicating the directional impact of a *Yes* answer. If a (+) item receives a *Yes* response, it increases blockchain technology's suitability for the financial service, while a *Yes* answer to a (-) item decreases its suitability. A (0) item's impact is not straightforward and can lean either way. These (0) items demand further discussion. For instance, *Item 2.1: Are multiple parties involved in the requirement for write access?* (+) presents two adjacent questions. If both questions - the involvement of multiple parties and the requirement for extensive write access - receive a *Yes* answer, the financial service indicates high applicability for blockchain technology. The sum of all positive items provides a final indication of blockchain technology's suitability for a specific financial service.

Table 4.1 The Assessment Framework

#	Item	Impact	Literature
1.1	Is the service significant enough to justify the costs of a decentralized network?	(+)	[1: 61]
1.2	Does a well-functioning solution already exist?	(-)	[42][4: 62]
1.3	Does the service require an unambiguous regulatory setup?	(-)	[108, 71][5: 49]
2.1	Are multiple parties involved in the requirement for write access?	(+)	[58, 94, 130, 78, 69, 42]
2.2	Do the parties have conflicts of trust or interest alignment?	(+)	[58, 68, 94, 130, 78]
2.3	Is a trusted third party feasible?	(-)	[58, 132, 94, 130, 42]
2.4	Does the service benefit from strict control and a majority-based approach for protocol updates?	(+)	[104, 93, 69, 42]
3.1	Is the service exposed to settlement risks?	(+/0/-)	[42, 63][6: 53]
3.2	Does the service store many nontransactional data in each transaction?	(0/-)	[132, 78, 69][3, 5: 122, 49]
3.3	Does the service require high throughput?	(0/-)	[78, 69, 42][4: 62]
3.4	Does the service benefit from blockchain technology's specific data properties, particularly immutability and transparency?	(+/0)	[78, 68, 94, 42][2, 4: 121, 62]
3.5	Does the service require privacy?	(+/0/-)	[108, 42, 31][1, 3: 61, 122]
3.6	Which identity level is required for the service?	(+/0/-)	[42, 31, 47][1, 3: 61, 122]
4.1	Does the service manage digital value exchange?	(+)	[130][1: 61]
4.2	Does the service require significant access to external information?	(-)	[1: 61]
4.3	Is the service exposed to increased counterparty risk?	(+)	[4: 62]
4.4	Does the service benefit from programmability and automation?	(+)	[60][2: 121]
4.5	Does the service require complex computation?	(-)	[5, 6: 49, 53]
4.6	Does the service benefit from open-source development?	(+)	[1, 5: 49, 61]
4.7	Does the service benefit from composability?	(+)	[2: 122][69]
4.8	Does the service benefit from an open-access environment?	(+)	[1, 4: 61, 62]
4.9	Does the service require centralized governance?	(-)	[120, 94]
4.10	Is the service suitable for game-theoretical mechanisms and strong market forces?	(+)	[63][3, 5: 122, 53]
4.11	Does the service require significant conflict resolution?	(-)	[108, 31]

4.4 Evaluation

In line with [116], a brief evaluation of the framework is applied to a combined financial service developed in [Paper 4: 62]. The evaluation seeks to exemplify how to apply the assessment framework to a financial service. Motivated by a large number of people in emerging economies without access to basic financial services [28, 7] and the high cost of banking services [51, 21], the paper explores the use of blockchain technology to create a prototype capable of performing three fundamental financial services: processing financial transactions, maintaining a savings account, and distributing targeted stimulus payments. To fund the stimulus payments, the prototype integrates its asset reserve into a public DeFi money market for yield generation. Thus, the paper examines the fascinating dilemma of a service that requires low-cost, scalable transactions while relying on blockchain’s trust guarantees to custody and govern the global asset reserve. As a result, the implemented prototype suggests building on two blockchains: one blockchain to account for the high trust requirements and DeFi integration, and another for performant microtransactions. Further, the paper analyzes the usage of the prototype through a network analysis nine months after deployment. The publication is particularly suited for the evaluation as it builds a blockchain prototype for emerging economies and subsequently analyzes its usage during a nine-month trial period, thus combining both leading research methodologies of the Ph.D. The evaluation is depicted in Table 4.2.

Evaluation	Score
1.1 The paper’s artifact explores blockchain technology’s capability to support a financial service that facilitates transactions at scale and targeted stimulus payments. The artifact enables users to execute financial transactions among one another and receive stimulus payments from a global asset reserve, directly to their digital wallets which function similarly to savings accounts. The importance of this financial infrastructure is highlighted especially in emerging economies, where approximately 1.7 billion people remain unbanked [29], a factor strongly tied to poverty [84]. A system for targeted stimulus payments could serve as an important tool, offering financial support in situations such as natural disasters or economic crises. <i>Yes</i> , the service is significant enough to justify the costs.	(+)
1.2 While there is financial infrastructure for saving accounts and payments, much room for improvement in these systems remains. Two key shortcomings are high tariffs and limited access. To our best knowledge, there was <i>no</i> widely adopted system specifically designed for targeted stimulus payments.	(+)

- 1.3 The regulation of financial services depends on the jurisdiction of individual countries, with each service requiring thorough and individual evaluation. Arguably, financial services are subject to regulation. A comprehensive analysis falls outside the scope of this thesis. (-)
-
- 2.1 Yes, the examined service involves multiple parties including banks, users, governments, and a non-governmental organization (NGO). (+)
- 2.2 Given the participation of multiple heterogenous parties, it can reasonably be assumed that they do not all share aligned interests. As suggested by the literature, these parties have yet to successfully implement a fully functional financial infrastructure [28, 7]. (+)
- 2.3 Although a trusted third party is generally a viable option, the parties involved have yet to successfully implement a fully functional financial infrastructure. To discuss if a third party is viable more regional and domain-specific knowledge is required, leading to a *neutral* evaluation. (0)
- 2.4 Updates to public blockchains are only implemented when there is a social consensus among all stakeholders, including infrastructure providers, users, and token holders. Given the participation of multiple heterogenous parties, it can reasonably be assumed that they are not all aligned in the current process. Thus, democratizing access to the base layer technology could prove beneficial. (+)
-
- 3.1 The three examined financial services are fairly simple and are characterized by micropayments, which neither have a high settlement risk in traditional finance nor a strong requirement for finality. However, latency may have a large role in micropayments as they are often used for day-to-day transactions at the point of sale. The distinct arguments lead to a *neutral* evaluation. (0)
- 3.2 The service records basic transactional information, including the sender, receiver, transaction currency, timestamp, and transaction value. The service does *not* require nontransactional data. (+)
- 3.3 Payment services, especially micropayments, require a high degree of throughput. (-)

- 3.4 The need for transparency, immutability, and auditability is key for targeted stimulus payments in emerging economies. The data-rich environment facilitates highly personalized stimulus payments. The study's artifact implements targeted incentive campaigns which are then analyzed using network analysis. (+)
- 3.5 All transactions of the artifact are publicly available on the blockchain. The users are pseudonymous; only the address is known. However, privacy is not complete, as sophisticated analysis tools could reveal a user. There is often a trade-off between users' need for privacy and the financial authorities' requirement for transparency in order to prevent fraudulent activities. This balance should be carefully considered when designing such a service. (0)
- 3.6 Blockchain technology's public key cryptography allows one to create multiple accounts. To avoid fraud with targeted stimulus payments the service benefits from authenticated identification of the stimulus payment receiver. In the artifact, this is partially implemented through the requirement of a signup e-mail address. (-)
-
- 4.1 The services are an immaterial exchange of value. (+)
- 4.2 The services do not require external information such as stock data or currency conversion rates. (+)
- 4.3 The three services involve multiple counterparties and intermediaries. Given the long-term nature of services such as maintaining savings accounts and custody of the global assets reserve, the risk from counterparties tends to escalate over time. Thus, this service is subject to heightened counterparty risk. (+)
- 4.4 The services benefit significantly from programmability and automation. High banking costs have been identified as a major obstacle in the traditional financial infrastructure for emerging economies [51, 21]. Autonomous smart contracts can be used to increase service offerings or to reduce the manual labor involved. (+)
- 4.5 The examined services are fairly simple, requiring only payment functions, account management, and the capability to process payments in bulk. *No* it does not require complex computation. (+)
- 4.6 The three services benefit from open-source development in multiple ways. Utilizing standards for tokens, transactions, and accounts not only enhance the system's development speed and security but also promote inclusivity by inviting additional entities to integrate with the system. (+)

- 4.7 The artifact relies on other financial services in the DeFi ecosystem, particularly in its management of the global asset reserve. The asset reserve generates income on a money market so as to pay for the stimulus payments. (+)
- 4.8 A major obstacle to the traditional financial system in emerging economies is the lack of access to basic financial services. The open accessibility of blockchain technology can provide significant value and potentially promotes financial inclusivity [7]. The large number of participants during the nine-month pilot phase is encouraging to reach global access with blockchain technology. The network analysis further indicated that the relationship has a temporary character with users interacting and then disengaging. (+)
- 4.9 The artifact's token holders govern the asset reserve, deciding on how to generate and distribute yield. Further, the configuration between both blockchains, one for performance and the other for high trust assumptions is governed by the token holders. Thus, the services depend on governance. (-)
- 4.10 The three services do not use sophisticated game-theoretical mechanisms. The services are in a natural equilibrium and are not prone to exploitation with the exception of the asset reserve. The asset reserve is integrated into an external DeFi service so as to generate yield. Owing to the financial integration it could be vulnerable to exploitation in both applications. However, both applications can be monitored in real-time, raising alerts in case of illicit behaviors. Further, the artifact implements very successful incentive campaigns, attracting users during a nine-month period. (+)
- 4.11 The three services are voluntary offers to users and typically involve microtransactions. This decreases the requirement for conflict resolution. Given that dispute or conflict resolution was deemed a low priority, no explicit process was implemented. Further, extensive testing was conducted prior to the deployment of the artifact so as to prevent technical errors in an immutable environment. Thus, conflict resolution does *not* have a significant role. (+)

Table 4.2: To exemplify a brief evaluation, the framework is applied to the financial service built in [Paper 4: 62]. Of 24 evaluation items, 16 received positive responses, while only four were negative. This positive result indicates that blockchain technology is an excellent fit with the examined financial services.

The assessment of the blockchain technology’s applicability to the three financial services implemented in the prototype indicates strong compatibility. Of 24 evaluation items, 16 received positive responses, while only four were negative. This positive result indicates that blockchain technology is an excellent fit with the examined financial services. Of the four negative items, throughput and identity were identified as significant challenges. The final iteration of the artifact responds to these challenges in two ways: First, it employs a customized blockchain for high-throughput payment transactions. This delegates the transactions to a blockchain optimized for low costs and high scalability. Further, the artifact leverages the public Ethereum blockchain for governance and the global asset reserve to ensure high security guarantees. Second, the artifact implements a whitelisting mechanism for accounts. To use the service, accounts must be associated with an e-mail address. As e-mail addresses are considered a weak form of identity [47], this measure helps to reduce the likelihood of multiple account creation but cannot totally circumvent it.

4.5 Discussion

The initial research phase was motivated by a desire to explore the technical conditions that guide DeFi applications on blockchain technology and to understand whether blockchain technology is generally suited for financial services. Thus, this essay first derives the fundamental properties of blockchain technology and the characteristics of the DeFi ecosystem with the help of existing literature, and one publication [Paper 1: 61]. The foundation chapter 2 concludes that there is a natural alignment between financial services and blockchain technology, primarily owing to their immateriality, the involvement of multiple unknown parties that require high trust, and a strong necessity for documentation [23, 51]. Blockchain’s interesting technical properties simultaneously enable the creation of novel financial services but are also restrictive in their capabilities. These properties create a distinct environment for financial services, one that is programmable, open, adversarial, and operates all the time [23, 25]. However, these preconditions may have distinct impacts on the broad range of financial services, further driven by regional differences [28, 7]. Blockchain is not a one-size-fits-all solution since, owing to the inherent trade-offs, each financial service requires specific consideration. This presents a challenging question for researchers, developers, and practitioners, hypothesizing about blockchain technology’s applicability to financial services.

The main part of this essay proposes a distilled framework that addresses the question of the extent to which blockchain technology is suitable for the implementation of financial services. The framework reflects the contemporary state of the art and has been informed by the six publications during the Ph.D.

program. Both blockchain technology and the principles that shape the DeFi application layer are subject to constant development. For instance, fairly recent blockchain layer updates have addressed key elements such as consensus and cryptographic signature schemes. Further, particularly experimentation with novel financial services in DeFi remains dynamic. Open DeFi has skyrocketed in popularity since 2020 owing to an explosive increase in the assets secured by the underlying blockchain technology [124][Paper 1: 61]. Thus, the framework is equally dynamic, potentially requiring future reassessments or even amendments.

Key research areas such as scalability, privacy, and identity are continuously explored. Considerable effort is being invested in the scalability of blockchain technology, with the expectation of increasing throughput [24, 77]. The transaction capacity may reach a level that would satisfy even the most performance-demanding of financial services. Recent iterations in blockchain technology are able to match the capabilities of VISA, one of the world's most significant centralized payment systems, with 24,000 Tps or 150 million transactions per day⁵. Although the comparison is not perfect, it offers a tangible benchmark. In light of these advancements, a reassessment of issues relating to system requirements such as [Item 3.2](#) and [Item 3.3](#) may be necessary. Simultaneously, ongoing research into privacy and identity may find a delicate balance between user privacy and preventing illicit behaviors. The emergence of novel privacy and identity techniques could lead to a reassessment of [Item 3.5](#) and [Item 3.6](#). Lastly, in a given timeframe clearer statements from regulatory authorities can be expected that can alter the assessment of regulatory requirements in [Item 1.3](#). Finally, new additions could be included, such as those that focus on a systemic perspective owing to the continuously advancing integration of DeFi applications [47][Paper 2: 121].

My Ph.D. process focused on popular public blockchain technology, exploring it from both a technical and an analytical perspective, while deliberately refraining from an in-depth exploration of three adjacent research areas: permissioned blockchain systems, comparisons to traditional financial infrastructure, and regulatory aspects of DeFi. Popular public blockchains such as the Ethereum blockchain have established strong network effects with a large number of developers and users, as well as applications [106, 23], resulting in a vibrant DeFi ecosystem that attracted billions of USDs in assets. In contrast to permissioned systems, public blockchains inherently foster composability, creating a competitive environment shaped by strong market forces with increasing systemic risks. The strong network effects and phenomena around composability and systemic risks were the main arguments to focus on public blockchain technology. It presents fascinating research opportunities that I explored during my Ph.D. process [Papers 2 and 6: 121, 53].

Assessing the applicability of financial services for blockchain technology

⁵<https://usa.visa.com/run-your-business/small-business-tools/retail.html>

benefits from comprehensive domain knowledge about the traditional financial infrastructure such as presented by [113]. This research area is closely tied to practice, as demonstrated by increasing publication activities by major financial institutions such as the Federal Reserve [106], the European Central Bank [17], and the Bank for International Settlement [6, 8]. The individual publications pursue this comparison only to a limited degree, for instance, [Paper 4: 62] compares existing mobile payment systems to the proposed blockchain solution, and [Paper 5: 49] uses interviews with industry experts for comparison. Despite the sophistication of existing financial infrastructure, its relevance is highly region-dependent. For instance, although VISA’s centralized payment systems perform very well, its coverage is geographically restricted [28, 7]. Financial services are not equally distributed globally, and blockchain technology may have bigger impacts in regions where the financial infrastructure is less developed [Paper 4: 62]. To fully understand blockchain technology’s potential for a specific financial service, extensive and geographically-diverse domain knowledge is necessary.

Overall, it is important to view the assessment framework as dynamic rather than static or exhaustive. The framework enables a balanced discussion among researchers and practitioners. It can be used as a starting point, emphasizing the research’s significance while identifying key questions and relevant literature. It also benefits strongly from additional financial domain knowledge and regulatory expertise.

Because two publications apply a prototyping research approach and three publications analyze DeFi phenomena with data, the framework is a culmination of a technical and analytical perspective. In pursuit of the research question, building and analyzing DeFi applications formed a synergetic relationship. The open-source ethos and public data enable insights from existing applications: analyzing the existing economic mechanisms and de facto uses assist financial services in their design decisions. The analysis of applications can inform the application’s design process prior to deployment by studying the design decisions and impact from already deployed applications [119]. After deployment, data can be used to monitor and assess the initial design, followed by updates for optimizations, an approach adopted in [Paper 4: 62]. This creates an intriguing evaluate-build-evaluate cycle that could lead to better applications and ultimately to a more resilient and attractive DeFi ecosystem.

A positive example of the relationship between analytics and development is the ongoing decentralization efforts of DeFi applications. Typically, DeFi applications target a distribution of governance power, aiming for a more resilient application [120, 10][Paper 1: 61]. This is commonly done through sophisticated governance mechanisms and token-based voting. As a pre-requirement, tokens must be distributed among key stakeholders. Analyzing existing protocols can provide valuable insights into the design of both, the

governance mechanism and the initial token distribution mechanism. A successful design can be constantly quantified by calculating the distribution of tokens [120][Paper 1: 61].

The relationship between building and analytics becomes increasingly important in an environment characterized by adversarial behaviors. Despite diligent development efforts, an application remains potentially susceptible to misuse [Paper 3: 122]. Hostile actors seeking profit can utilize an application in ways contrary to its intended design [139]. Given that an application is often designed without interdependency in mind, vulnerabilities can arise at the interface between two applications. The granularity and real-time properties of blockchain data enable sophisticated monitoring of individual applications and the entire DeFi ecosystem. The assets and liabilities of any application are publicly available and enable continual monitoring of an application, raising alerts in case of misuse [Paper 3: 122]. The same data-rich environment facilitates systemic research and market research, since all interactions with applications are accessible, regardless of whether the interaction was triggered by an EoA or another smart contract [65][Paper 2: 121].

Further abstracting away from the framework, allows a discussion of the benefits and challenges of DeFi. While utilizing the strengths of blockchain-based finance offers compelling potentials it is important to address the limitations. First, DeFi presents an opportunity for trust-minimized disintermediation. The formalized protocol of a blockchain, operated by incentivized participants, reduces the need for intermediaries, which may lead to significant cost savings and efficiency gains [106, 59]. This trust-minimized property represents a fundamental shift in how financial transactions are conducted. The democratization of access is another opportunity afforded by blockchain technology. The open access of blockchain-based platforms allows anyone, regardless of their geographical location or social status, to access and participate in these financial systems. This opens possibilities for financial inclusion and broader participation [106, 51].

DeFi's composability enables the synergetic interactions of applications and, together with the expansive design space, forms a vibrant ecosystem of financial services. DeFi applications can seamlessly integrate with and build on other applications, creating novel, ever-expanding financial services [Paper 2: 121]. This leads to a multi-component financial system [47] that functions as a single, coherent market [109]. Applications not only share computational resources but also users, who can easily move between competing protocols. Users benefit from this interoperability since they can seamlessly move their assets between different applications, thereby finding customized financial services that cater to their preferences.

Further, DeFi is built on economic incentives since each application aims to strike a balance to incentivize usage and maintain equilibrium. These sophisticated game-theoretical mechanisms promote competition and present

arbitrage opportunities. Opportunities for arbitrage arise when price discrepancies occur between different applications, allowing arbitrageurs to simultaneously buy and sell assets across platforms, profiting from the spread [Paper 6: 53]. In the process, they contribute to the overall efficiency of the DeFi ecosystem by maintaining price consistency and reducing discrepancies. An interoperable and competitive DeFi ecosystem shares liquidity and optimizes the distribution of assets under strong market forces [51].

Lastly, the data-rich environment of DeFi results in the largest dataset of a public financial system [109]. The valuable data enable sophisticated analysis benefitting market research, real-time monitoring of applications, and fraud detection [119][Paper 3: 122]. The system's financial integration can be analyzed at high granularity, which is beneficial for ongoing risk management of both an application and the financial system [65][Paper 2: 121].

However, despite leveraging the benefits and making deliberate design choices for the implementation of a financial service, blockchain-based financial applications still face challenges across every layer of the technology stack.

A significant challenge is the inherent costs of operating in a decentralized system, leading to performance limitations as well as relatively high computation and storage costs. Processing capabilities on the popular public blockchains are lower than traditional, centralized financial technology, which limits the performance and scalability of DeFi applications [70, 133]. Also, storage on a blockchain is expensive and the extent of smart contract execution is limited [5]. These current limitations restrict the range of functions and features that can be implemented in DeFi applications, especially considering that financial use cases often require high-frequency, low-latency technology. Addressing these scaling issues is crucial for the future of DeFi.

Another challenge is the technical security risk inherent when working in a nascent environment. DeFi applications are being developed on a fairly nascent technology while managing significant asset values. The rapid pace of development introduces the risk of unforeseen or hostile security breaches that can potentially result in substantial losses. Security breaches can be introduced at both the base layer [25, 63, 100] and the application layer [47, 96, 66]. This risk is magnified since the application's code is public, hostile actors are pseudonymous, and transactions are irreversible.

Further, the advanced integration of DeFi applications creates multiple interfaces that can present potential vulnerabilities. Vulnerabilities in one application can propagate and compromise other integrated applications, leading to a cascading effect through the entire ecosystem. For instance, assets are reused and accounted for on multiple applications in pursuit of higher yields, akin to rehypothecation in traditional finance [Paper 2: 121]. A vulnerability in one application can render the assets worthless across all integrated applications. Advanced integration also leads to additional layers of complexity, which extend to the data. Although DeFi platforms are inherently transpar-

ent, understanding and parsing the data often requires specialized knowledge, making DeFi analytics a challenging endeavor [39].

Another limitation of DeFi is the potential for illicit activities on applications owing to pseudonymous addresses [106]. Hostile actors can exploit the intended economic design of applications by disrupting their intended equilibrium [139]. For instance, [Paper 3: 122] explores and quantifies how NFTs are traded to signal higher prices or avoid taxes, representing a misuse of DeFi's NFT marketplaces. While a blockchain's pseudonymity offers privacy advantages, it also introduces the possibility of misuse.

In sum, financial services are particularly suited for blockchain technology if existing systems present significant opportunities for improvement. For instance, [Paper 4: 62] explores blockchain services for emerging economies, where the regional financial infrastructure is dysfunctional or is subject to high costs. Blockchain technology is potentially beneficial if the financial service at hand involves untrusted participants. In a sufficiently decentralized blockchain-based financial service, no single entity has full control, making it more difficult for individual participants to manipulate the systems. The inherent transparency allows all participants to audit the network. Further, blockchain technology is highly applicable if a financial service can effectively maximize DeFi's potential. The design space for financial services on a blockchain is vast [Paper 5: 49], enabled by the programmability of smart contracts and the exponential possibilities to integrate into other applications. In the pursuit of high yields or customized services, users can seamlessly move assets between the different applications. The resulting financial ecosystem is simultaneously competitive and collaborative and is therefore shaped by strong market forces [51].

The highlighted limitations present intriguing open questions for DeFi research. Although composability is an interesting property, many applications are not designed with interdependency in mind, leading to vulnerabilities at the interface between two applications. The trend toward increasingly complex applications and advanced integration highlights the importance of resilient application design that can otherwise lead to a systemic crisis [47]. Systemic risk studies that analyze the underlying integrations of traditional financial systems have become increasingly popular so as to gain a better understanding of these risks in DeFi [109, 107].

An interesting yet underdeveloped research area is in the integration between traditional finance and DeFi. Currently, DeFi services are predominantly implemented by entities originating within the DeFi ecosystem, while established companies are largely absent. Such absence is notable in capital-intensive industries such as insurance, potentially a reason for the relative underdevelopment of insurance services in DeFi. The merger of traditional and decentralized finance could balance the regulatory and institutional trust of established companies with the technological advancements of DeFi. How-

ever, the existing technological divergences between the two systems and the regulatory challenges associated with this integration make it an intriguing and complex subject.

Conclusion

This essay proposes guiding principles that determine the extent to which blockchain technology can be used for the implementation of specific financial services. The essay starts by deriving the fundamental properties of blockchain technology as well as the characteristics and values of DeFi, aiming to understand the technical conditions that guide financial services on blockchain technology. This initial chapter concludes that there is a natural alignment between financial services and blockchain technology, primarily owing to their immateriality, the involvement of multiple unknown parties requiring a strong trust, and a strong necessity for documentation [23, 51]. However, blockchain technology is not necessarily required for every financial service. Blockchain technology can be restrictive in its technical capabilities and creates a distinct environment for financial services, one that is programmable, openly accessible, adversarial, and continuously operational [23, 25]. These preconditions have distinct impacts on the broad range of financial services, further driven by regional differences [28, 7]. Thus, financial services benefit from an individual assessment of blockchain technology’s applicability.

Following this call, the essay culminates in an assessment framework that contributes to the understanding of the extent to which blockchain technology is suitable for the implementation of financial services. The framework consists of 24 evaluation items, that are drawn from the generalizable insights from six academic publications and additional literature.

In sum, financial services are particularly suited for blockchain technology if existing systems present significant opportunities for improvement. For instance, [Paper 4: 62] explores blockchain services for emerging economies, where the regional financial infrastructure is dysfunctional or is subject to high costs. Blockchain technology is potentially beneficial if the financial service at hand involves untrusted participants. In a sufficiently decentralized blockchain-based financial service, no single entity has full control, making it more difficult for individual participants to manipulate the systems. The inherent transparency allows all participants to audit the network. Further, blockchain technology is highly applicable if a financial service can effectively maximize DeFi’s potential. The design space for financial services on a blockchain is vast [Paper 5: 49], enabled by the programmability of smart

contracts and the exponential possibilities to integrate into other applications. In the pursuit of high yields or customized services, users can seamlessly move assets between the different applications. The resulting financial ecosystem is simultaneously competitive and collaborative and is therefore shaped by strong market forces [51].

The individual publications form the central part of the Ph.D. program and share the overarching domain of DeFi. The publications follow a research approach based on prototyping or network analysis. The prototypes built in the publications serve as proofs of concept and aid in understanding the problem. The analytical work explores the de facto use of existing applications, aiming to broaden the discussion from purely technical to encompass human interactions and economics. Thus, the publications are representative in that the method of building and analyzing DeFi applications forms a synergetic relationship. The analysis of the available data can inform the application design process prior to deployment, while post-deployment data can be used to monitor, assess, and update the initial design for optimizations. This facilitates objective discussions and creates an evaluate-build-evaluate cycle that may lead to better applications and ultimately to a more resilient and attractive DeFi ecosystem. Blockchain technology's transparency also helps the broader DeFi ecosystem to monitor financial integration and systemic risks.

This thesis is a cumulative dissertation and concludes my Ph.D. program. It consists of this essay and a collection of six publications, each published during my Ph.D. program. It contributes to the existing knowledge, particularly in the IS discipline, an interdisciplinary field that combines elements of computer science, economics, and social science [9, 105]. Being a blend of distributed systems, cryptography, and economic incentives [70], blockchain technology is well-suited for an interdisciplinary research approach [112, 90]. To the best knowledge, this assessment framework is the first structured approach to objectify blockchain technology's applicability to financial services. The framework seeks to provide researchers and practitioners with a valuable tool, while simultaneously encouraging further research opportunities. Further, the individual publications make unique contributions to their respective fields. Specifically, the publications propose blockchain-based building blocks for financial services in the areas of targeted payments, asset reserves, equity, and crowdfunding. The empirical publications propose algorithms, quantifying domain-specific phenomena such as system integration, illicit behaviors in marketplaces, and arbitrage.

With the impressive growth of DeFi, amassing billions of USDs in assets, the significance of this research continues to evolve, potentially addressing critical questions toward establishing a more resilient financial system based on blockchain technology.

Despite the extensive effort, my Ph.D. process and the culminating the-

sis has limitations. As the assessment framework is discussed with relevant criticism, the framework should be considered to be dynamic, generalizing, and not exhaustive. Further, the framework's scoring system does not offer definitive answers; instead, it serves to encourage balanced discussion among researchers and practitioners. It can serve as a starting point, emphasizing the research's significance while identifying key questions and relevant literature. Further, A focus on public DeFi inevitably led to certain compromises, such as limited attention to regulatory assessments and comparison to traditional financial systems [98].

Discussing the assessment framework, this essay raises two intriguing avenues for future research. First, DeFi tends toward higher financial integration levels and increasingly complex applications [47][Paper 2: 121]. This encourages a systemic perspective to understand systemic risks inherent in a system and is akin to disciplines in traditional financial research. Given that an application is often designed without interdependency in mind, vulnerabilities can arise in the interface between two applications. Second, the integration between traditional finance and DeFi represents an interesting yet underexplored research area. Currently, DeFi services are predominantly provided by companies originating within the DeFi ecosystem, with established companies from capital-intensive industries such as insurance notably absent. The merger of traditional finance and DeFi could balance between the regulatory and institutional trust inherent in traditional finance with DeFi's technological potential.

In conclusion, while blockchain technology and DeFi may not be suitable for every single financial service, they present strong potentials to rethink the understanding of and interaction with the existing financial infrastructure. This technology could facilitate a more equitable and more efficient financial ecosystem, with intriguing experimentation already underway. However, the achievement of such potential success is a long-term process. It is crucial to acknowledge that blockchain technology still requires ongoing research and development in key areas such as scalability, regulation, and identity.

References

- [1] Altay Aksulu and Michael Wade. A Comprehensive Review and Synthesis of Open Source Research. *Journal of the Association for Information Systems*, 11(11):576–656, 2010.
- [2] Hamda Al-Breiki, Muhammad Habib Ur Rehman, Khaled Salah, and Davor Svetinovic. Trustworthy Blockchain Oracles: Review, Comparison, and Open Research Challenges. *IEEE Access*, 8, 2020.
- [3] Emmanuelle Anceaume, Antonella Pozzo, Thibault Rieutord, and Sara Tucci-Piergiovanni. On Finality in Blockchains. *arXiv*, 2020. URL <http://arxiv.org/abs/2012.10172>.
- [4] Andreas Antonopoulos. *Mastering Bitcoin*. O’Reilly Media, Sebastopol, California, USA, 2014.
- [5] Andreas Antonopoulos and Gavin Wood. *Mastering Ethereum*. O’Reilly Media, Sebastopol, California, USA, 2018.
- [6] Sirio Aramonte, Wenqian Huang, and Andreas Schrimpf. DeFi Risks and the Decentralisation Illusion. *BIS Quarterly Review*, 2021.
- [7] Oya Pinar Ardic, Maximilien Heimann, and Nataliya Mylenko. Access to Financial Services and the Financial Inclusion Agenda around the World: A cross-country Analysis with a new Data Set. *Worldbank Policy Research Working Paper*, 2011.
- [8] Raphael Auer, Bernhard Haslhofer, Stefan Kitzler, Pietro Saggese, and Friedhelm Victor. The Technology of Decentralized Finance (DeFi). *BIS Working Papers 1066, Bank for International Settlements*, 2023.
- [9] Chrisanthi Avgerou. Information Systems: What Sort of Science is it? *Omega*, 28:567–579, 2000.
- [10] Henrik Axelsen, Johannes Rude Jensen, and Omri Ross. When is a DAO Decentralized? *Complex Systems Informatics and Modeling Quarterly*, 31(4):51–75, 2022.

- [11] Marco Bardoscia, Paolo Barucca, Stefano Battiston, Fabio Caccioli, Giulio Cimini, Diego Garlaschelli, Fabio Saracco, Tiziano Squartini, and Guido Caldarelli. The Physics of Financial Networks. *Nature Reviews Physics*, 3(7):490–507, 2021.
- [12] Massimo Bartoletti and Livio Pompianu. An Empirical Analysis of Smart Contracts: Platforms, Applications, and Design Patterns. In *International Conference on Financial Cryptography and Data Security (FC 2017)*, pages 494–509, 2017.
- [13] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch-Lafuente. A Theory of Automated Market Makers in DeFi. In *International Conference on Coordination Languages and Models*, pages 168–187, 2021.
- [14] Roman Beck, Jacob Stenum Czepluch, Nikolaj Lollike, and Simon Malone. Blockchain - The Gateway to Trustfree Transactions. In *European Conference on Information Systems (ECIS 2016)*, 2016.
- [15] Roman Beck, Christoph Müller-Bloch, and John King. Governance in the Blockchain Economy: A Framework and Research Agenda. *Journal of the Association for Information Systems*, 19(10), 2018.
- [16] Balázs Bodó and Primavera De Filippi. Trust in Context: The Impact of Regulation on Blockchain and DeFi. *Regulation and Governance*, 2022.
- [17] Alexandra Born, Isabella Gschossmann, Alexander Hodbod, Claudia Lambert, and Antonella Pellicani. Decentralised Finance – a new unregulated non-bank System? *ECB*, 2022. URL https://www.ecb.europa.eu/pub/financial-stability/macprudential-bulletin/focus/2022/html/ecb.mpbu202207_focus1.en.html.
- [18] Antonio Briola, David Vidal-Tomás, Yuanrong Wang, and Tomaso Aste. Anatomy of a Stablecoin’s Failure: The Terra-Luna Case. *Finance Research Letters*, 51, 2023.
- [19] Vitalik Buterin. A Next Generation Smart Contract and Decentralized Application Protocol. *Whitepaper*, 2013. URL https://ethereum.org/669c9e2e2027310b6b3cdce6e1c52962/Ethereum_Whitepaper_-_Buterin_2014.pdf.
- [20] Fabio Castiglionesi, Fabio Feriozzi, and Guido Lorenzoni. Financial Integration and Liquidity Crises. *Management Science*, 65(3):955–975, 2017.
- [21] Bhaskar Chakravorti. The hidden Costs of Cash. *Harvard Business Review*, 18, 2015.

- [22] Weili Chen, Tuo Zhang, Zhiguang Chen, Zibin Zheng, and Yutong Lu. Traveling the Token World: A Graph Analysis of Ethereum ERC20 Token Ecosystem. In *The Web Conference 2020 (WWW 2020)*, pages 1411–1421, 2020.
- [23] Yan Chen and Cristiano Bellavitis. Blockchain Disruption and Decentralized Finance: The Rise of Decentralized Business Models. *Journal of Business Venturing Insights*, 13, 2020.
- [24] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. On Scaling Decentralized Blockchains. In *International Conference on Financial Cryptography and Data Security (FC 2016)*, pages 106–125, 2016.
- [25] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. In *IEEE Symposium on Security and Privacy*, pages 910–927, 2020.
- [26] Primavera De Filippi and Samer Hassan. Blockchain Technology as a regulatory Technology: From Code is Law to Law is Code. *First Monday*, 21(12), 2016.
- [27] Primavera De Filippi, Morshed Mannan, and Wessel Reijers. Blockchain as a Confidence Machine: The Problem of Trust & Challenges of Governance. *Technology in Society*, 62, 2020.
- [28] Asli Demirgüç-Kunt and Leora Klapper. Measuring Financial Inclusion: Explaining Variation in Use of Financial Services across and within Countries. *Brookings Papers on Economic Activity*, 2013(1):279–340, 2013.
- [29] Asli Demirgüç-Kunt, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess. *The Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution*. World Bank Reports, 2018.
- [30] Anil Donmez and Alexander Karaivanov. Transaction Fee Economics in the Ethereum Blockchain. *Economic Inquiry*, 60(1):265–292, 2022.
- [31] Benjamin Egelund-Müller, Martin Elsman, Fritz Henglein, and Omri Ross. Automated Execution of Financial Contracts on Blockchains. *Business and Information Systems Engineering*, 59(6):457–467, 2017.
- [32] Shayan Eskandari, Mehdi Salehi, Wanyun Catherine Gu, and Jeremy Clark. SoK: Oracles from the Ground Truth to Market Manipulation. In *AFT 2021 - Proceedings of the 2021 3rd ACM Conference on Advances in Financial Technologies*, pages 127–141, 2021.

- [33] Benedict Faber, Georg Cappelen Michelet, Niklas Weidmann, Raghava Rao Mukkamala, and Ravi Vatrappu. BPDIMS: A Blockchain-based Personal Data and Identity Management System. *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 6:6855–6864, 2019.
- [34] B. A. Farbey. Structural Models: An Introduction to the Theory of Directed Graphs. *Journal of the Operational Research Society*, 17(2):202–203, 1966.
- [35] Qi Feng, Debiao He, Sherali Zeadally, Muhammad Khurram Khan, and Neeraj Kumar. A Survey on Privacy Protection in Blockchain System. *Journal of Network and Computer Applications*, 126:45–58, 2019.
- [36] Michael Fröwis, Andreas Fuchs, and Rainer Böhme. Detecting Token Systems on Ethereum. In *International Conference on Financial Cryptography and Data Security (FC 2019)*, pages 93–112, 2019.
- [37] W. Gawlikowicz, B. Mannerings, T. Rudolph, and D. Šiška. Market Based Mechanisms for Incentivising Exchange Liquidity Provision. In *International Conference on Financial Cryptography and Data Security (FC21)*, pages 80–96, 2021.
- [38] George M. Giaglis and Kalliopi N. Kypriotaki. Towards an Agenda for Information Systems Research on Digital Currencies and Bitcoin. In *International Conference on Business Information Systems*, pages 3–13. Springer Verlag, 2014.
- [39] Florian Glaser. Pervasive Decentralisation of Digital Infrastructures: A Framework for Blockchain enabled System and Use Case Analysis. In *Hawaii International Conference on System Sciences (HICSS 2017)*, 2017.
- [40] Florian Glaser and Luis Bezenberger. Beyond Cryptocurrencies - A Taxonomy of Decentralized Consensus Systems. In *European Conference on Information Systems (ECIS 2015)*, 2015.
- [41] Paul Glasserman and H Peyton Young. Contagion in Financial Networks. *Journal of Economic Literature*, 54(3), 2016.
- [42] Sri Nikhil Gupta Gourisetti, Michael Mylrea, and HIRAK Patangia. Evaluation and Demonstration of Blockchain Applicability Framework. *IEEE Transactions on Engineering Management*, 67(4):1142–1156, 2020.
- [43] Vincent Gramlich, Tobias Guggenberger, Marc Principato, Benjamin Schellinger, and Nils Urbach. A Multivocal Literature Review of Decentralized Finance: Current Knowledge and Future Research Avenues. *Electronic Markets*, 33(1):11, 2023.

- [44] Laura Grassi, Davide Lanfranchi, Alessandro Faes, and Filippo Maria Renga. Do we still need Financial Intermediation? The Case of Decentralized Finance – DeFi. *Qualitative Research in Accounting & Management*, 19(3):323–347, 2022.
- [45] Shirley Gregor and Alan R Hevner. Positioning and Presenting Design Science Research with Maximum Impact. *MIS Quarterly*, 37(2):337–355, 2013.
- [46] J Grigo, P Hansen, A Patz, and V von Wachter. Decentralized Finance (DeFi) - A new Fintech Revolution. *Bitkom*, 2020. URL https://www.bitkom.org/sites/main/files/2020-07/200729_whitepaper_decentralized-finance.pdf.
- [47] Lewis Gudgeon, Daniel Perez, Dominik Harz, Benjamin Livshits, and Arthur Gervais. The Decentralized Financial Crisis. In *IEEE Crypto Valley Conference on Blockchain Technology*, 2020.
- [48] Lewis Gudgeon, Sam Werner, Daniel Perez, and William J. Knottenbelt. DeFi Protocols for Loanable Funds: Interest Rates, Liquidity and Market Efficiency. In *2nd ACM Conference on Advances in Financial Technologies (AFT2020)*, pages 92–112, 2020.
- [49] Tobias Guggenberger, Benjamin Schellinger, Victor von Wachter, and Nils Urbach. Kickstarting Blockchain: Designing Blockchain-based Tokens for Equity Crowdfunding. *Electronic Commerce Research*, pages 1–35, 2023.
- [50] Harry Halpin and Marta Piekarska. Introduction to Security and Privacy on the Blockchain. In *2017 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 1–3, 2017.
- [51] Campbell Harvey, Ashwin Ramachandran, and Joey Santoro. *DeFi and the Future of Finance*. Wiley, 2021.
- [52] Florian Hawlitschek, Benedikt Notheisen, and Timm Teubner. The Limits of trust-free Systems: A Literature Review on Blockchain Technology and Trust in the sharing Economy. *Electronic Commerce Research and Applications*, 29:50–63, 2018.
- [53] Songrun He, Asaf Manela, Omri Ross, and Victor von Wachter. Fundamentals of Perpetual Futures. *SSRN Electronic Journal*, 2022. URL <https://papers.ssrn.com/abstract=4301150>.
- [54] Christine V Helliard, Louise Crawford, Laura Rocca, Claudio Teodori, and Monica Veneziani. Permissionless and Permissioned blockchain diffusion. *International Journal of Information Management*, 54, 2020.

- [55] Fritz Henglein. Blockchain Deconstructed. In *Abstract from 3rd Symposium on Distributed Ledger Technology (DLT3)*, 2018.
- [56] Alan Hevner and Samir Chatterjee. *Design Research in Information Systems*, volume 22 of *Integrated Series in Information Systems*. Springer US, Boston, Massachusetts, USA, 2010.
- [57] Alan R Hevner, Salvatore T March, Jinsoo Park, and Sudha Ram. Design Science in Information Systems Research. *MIS Quarterly*, 28(1):75–105, 2004.
- [58] Jens J Hunhevicz and Daniel M Hall. Do you need a Blockchain in Construction? Use Case Categories and Decision Framework for DLT Design Options. *Advanced Engineering Informatics*, 45, 2020.
- [59] Johannes Rude Jensen and Omri Ross. Settlement with Distributed Ledger Technology. In *International Conference on Information Systems (ICIS 2020)*, 2020.
- [60] Johannes Rude Jensen, Omri Ross, and Victor von Wachter. Leveraged Trading on Blockchain Technology. *arXiv*, 2020. URL <https://arxiv.org/abs/2102.13488>.
- [61] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. An Introduction to Decentralized Finance (DeFi). *Complex Systems Informatics and Modeling Quarterly*, 26(3):46–54, 2021.
- [62] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. Blockchain-based Financial Infrastructure for Emerging Economies. In *European Conference for Information Systems (ECIS 2022)*, 2022.
- [63] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. Multi-block MEV. *arXiv*, 2023. URL <http://arxiv.org/abs/2303.04430>.
- [64] Changhoon Kang, Chaehyeon Lee, Kyungchan Ko, Jongsoo Woo, and James Won-Ki Hong. De-Anonymization of the Bitcoin Network Using Address Clustering. In *International Conference on Blockchain and Trustworthy Systems*, pages 489–501, 2020.
- [65] Stefan Kitzler, Friedhelm Victor, Pietro Saggese, and Bernhard Haslhofer. Disentangling Decentralized Finance Compositions. *ACM Transactions on the Web*, 17(2), 2023.
- [66] Aariah Klages-Mundt and Andreea Minca. (In)Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks. *Cryptoeconomic Systems*, 1(2), 2021.

- [67] Aariah Klages-Mundt, Dominik Harz, Lewis Gudgeon, Jun-You Liu, and Andreea Minca. Stablecoins 2.0: Economic Foundations and Risk-based Models. In *ACM Conference on Advances in Financial Technologies*, pages 59–79, New York, NY, USA, 2020. ACM.
- [68] Sandra Klein, Wolfgang Prinz, and Wolfgang Gräther. A Use Case Identification Framework and Use Case Canvas for Identifying and Exploring relevant Blockchain Opportunities. In *Proceedings of the 1st ERCIM Blockchain Workshop 2018, European Society for Socially Embedded Technologies*, 2018.
- [69] Tommy Koens and Erik Poll. What Blockchain Alternative Do You Need? In *Conference on Data Privacy Management, Cryptocurrencies and Blockchain Technology (DTM and CTB 2018)*, pages 113–129, 2018.
- [70] John Kolb, Moustafa AbdelBaky, Randy H Katz, and David E Culler. Core Concepts, Challenges, and Future Directions in Blockchain. *ACM Computing Surveys*, 53(1):1–39, 2020.
- [71] Olga Labazova. Towards a Framework for Evaluation of Blockchain Implementations. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [72] Olga Labazova, Tobias Dehling, and Ali Sunyaev. From Hype to Reality: A Taxonomy of Blockchain Applications. In *Hawaii International Conference on System Sciences (HICS 2019)*, pages 4555–4564, 2019.
- [73] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.
- [74] Jei Young Lee. A Decentralized Token Economy: How Blockchain and Cryptocurrency can revolutionize Business. *Business Horizons*, 62(6):773–784, 2019.
- [75] Juho Lindman, Virpi Kristiina Tuunainen, and Matti Rossi. Opportunities and Risks of Blockchain Technologies: A Research Agenda. In *Hawaii International Conference on System Sciences (HICSS 2017)*, 2017.
- [76] Yang Liu, Debiao He, Mohammad S. Obaidat, Neeraj Kumar, Muhammad Khurram Khan, and Kim-Kwang Raymond Choo. Blockchain-based Identity Management Systems: A Review. *Journal of Network and Computer Applications*, 166, 2020.
- [77] Yizhong Liu, Jianwei Liu, Marcos Antonio Vaz Salles, Zongyang Zhang, Tong Li, Bin Hu, Fritz Henglein, and Rongxing Lu. Building Blocks of Sharding Blockchain Systems: Concepts, Approaches, and open Problems. *Computer Science Review*, 46, 2022.

- [78] Sin Kuang Lo, Xiwei Xu, Yin Kia Chiam, and Qinghua Lu. Evaluating Suitability of Applying Blockchain. In *2017 22nd International Conference on Engineering of Complex Computer Systems (ICECCS)*, pages 158–161, 2017.
- [79] Fabrice Lumineau, Wenqian Wang, and Oliver Schilke. Blockchain Governance—A New Way of Organizing Collaborations? *Organization Science*, 32(2):500–521, 2021.
- [80] Loi Luu, Duc Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making Smart Contracts Smarter. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 254–269, 2016.
- [81] Salvatore T. March and Gerald F. Smith. Design and Natural Science Research on Information Technology. *Decision Support Systems*, 15(4): 251–266, 1995.
- [82] Eva Meyer, Isabell Welpel, and Philipp Sandner. Decentralized Finance - A Systematic Literature Review and Research Direction. In *European Conference on Information Systems (ECIS 2022)*, 2022.
- [83] David Mills, Kathy Wang, Brendan Malone, Anjana Ravi, Jeff Marquardt, Clinton Chen, Anton Badev, Timothy Brezinski, Linda Fahy, Kimberley Liao, Vanessa Kargenian, Max Ellithorpe, Wendy Ng, and Maria Baird. Distributed Ledger Technology in Payments, Clearing, and Settlement. *Finance and Economics Discussion Series*, 2016(95), 2016.
- [84] Lakshmi Mohan and Devendra Potnis. Real-time decision-making to serve the Unbanked Poor in the Developing World. *SIGMIS-CPR 2017 - Proceedings of the 2017 ACM SIGMIS Conference on Computers and People Research*, pages 183–184, 2017.
- [85] Amani Moin, Kevin Sekniqi, and Emin Gun Sirer. SoK: A Classification Framework for Stablecoin Designs. In *International Conference on Financial Cryptography and Data Security (FC 2020)*, 2020.
- [86] Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. Mapping the NFT Revolution: Market Trends, Trade Networks, and Visual Features. *Scientific Reports*, 11(1), 2021.
- [87] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. *Whitepaper*, 2008. URL www.bitcoin.org.
- [88] Mark Newman. *Networks: An Introduction*. Oxford University Press, Oxford, UK, 2010.

- [89] Nick Szabo. First Monday: Formalizing and Securing Relationships on Public Networks. *First Monday*, 2(9), 1997.
- [90] Benedikt Notheisen, Florian Hawlitschek, and Christof Weinhardt. Breaking down the Blockchain Hype - Towards a Blockchain Market Engineering Approach. In *European Conference on Information Systems (ECIS 2017)*, pages 1062–1080, 2017.
- [91] OECD. Why Decentralised Finance (DeFi) Matters and the Policy Implications. *Position Paper*, 2022. URL <https://www.oecd.org/daf/fin/financial-markets/Why-Decentralised-Finance-DeFi-Matters-and-the-Policy-Implications.pdf>.
- [92] José Parra Moyano and Omri Ross. KYC Optimization Using Distributed Ledger Technology. *Business and Information Systems Engineering*, 59(6): 411–423, 2017.
- [93] Morgen E. Peck. Blockchain World - Do you need a Blockchain? This Chart will tell you if the Technology can solve your Problem. *IEEE Spectrum*, 54(10):38–60, 2017.
- [94] Asger B. Pedersen, Marten Risius, and Roman Beck. A ten-step Decision Path to Determine When to Use Blockchain Technologies. *MIS Quarterly Executive*, 18(2):99–115, 2019.
- [95] Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, 24(3):45–77, 2007.
- [96] Daniel Perez, Sam M Werner, Jiahua Xu, and Benjamin Livshits. Liquidations: DeFi on a Knife-edge. In *International Conference on Financial Cryptography and Data Security (FC 2021)*, 2021.
- [97] Julien Polge, Jérémy Robert, and Yves Le Traon. Permissioned Blockchain Frameworks in the Industry: A Comparison. *ICT Express*, 7(2):229–233, 2021.
- [98] Kaihua Qin, Liyi Zhou, Yaroslav Afonin, Ludovico Lazzaretti, and Arthur Gervais. CeFi vs. DeFi - Comparing Centralized to Decentralized Finance. *arXiv*, 2021. URL <http://arxiv.org/abs/2106.08157>.
- [99] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. Attacking the DeFi Ecosystem with Flash Loans for Fun and Profit. In *International Conference on Financial Cryptography and Data Security (FC 2021)*, pages 3–32, 2021.

- [100] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying Blockchain Extractable Value: How dark is the Forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214, 2022.
- [101] Tripti Rathee and Parvinder Singh. A Systematic Literature Mapping on Secure Identity Management using Blockchain Technology. *Journal of King Saud University - Computer and Information Sciences*, 34(8):5782–5796, 2022.
- [102] Ferdinand Regner, André Schweizer, and Nils Urbach. NFTs in Practice - Non-fungible Tokens as core component of a Blockchain-based Event Ticketing Application. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [103] Marten Risius and Kai Spohrer. A Blockchain Research Framework. *Business & Information Systems Engineering*, 59(6):385–409, 2017.
- [104] Omri Ross, Johannes Rude Jensen, and Truls Asheim. Assets under Tokenization: Can Blockchain Technology Improve Post-Trade Processing? In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [105] Suprateek Sarker, Sutirtha Chatterjee, Xiao Xiao, and Amany Elbanna. The Sociotechnical Axis of Cohesion for the IS Discipline: Its Historical Legacy and its Continued Relevance. *MIS Quarterly*, 43(3):695–719, 2019.
- [106] Fabian Schär. Decentralized Finance: On Blockchain- and Smart Contract-based Financial Markets. *St. Louis Reserve Bank Review*, 103(2), 2021.
- [107] Frank Schweitzer, Giorgio Fagiolo, Didier Sornette, Fernando Vega-Redondo, Alessandro Vespignani, and Douglas R. White. Economic Networks: The New Challenges. *Science*, 325(5939):422–425, 2009.
- [108] Brian A. Scriber. A Framework for Determining Blockchain Applicability. *IEEE Software*, 35(4):70–77, 2018.
- [109] Shahar Somin, Yaniv Altshuler, Goren Gordon, Alex Sandy 'Pentland', and Erez Shmueli. Network Dynamics of a Financial Ecosystem. *Scientific Reports*, 10(1), 2020.
- [110] Palina Tolmach, Yi Li, Shang Wei Lin, and Yang Liu. Formal Analysis of Composable DeFi Protocols. In *International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.*, pages 149–161, 2021.
- [111] Stefan Tönnissen and Frank Teuteberg. Towards a Taxonomy for Smart Contracts. In *European Conference on Information Systems (ECIS 2018)*, 2018.

- [112] Horst Treiblmaier, Melanie Swan, Primavera De Filippi, Mary Lacity, Thomas Hardjono, and Henry Kim. What’s Next in Blockchain Research? An Identification of Key Topics Using a Multidisciplinary Perspective. *Data Base for Advances in Information Systems*, 52(1):27–52, 2021.
- [113] Philip Treleaven, Richard Gendal Brown, and Danny Yang. Blockchain Technology in Finance. *Computer*, 50(9):14–17, 2017.
- [114] Florian Tschorsch and Bjorn Scheuermann. Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. *IEEE Communications Surveys & Tutorials*, 18(3):2084–2123, 2016.
- [115] Gerry Tsoukalas and Brett Hemenway Falk. Token-Weighted Crowdsourcing. *Management Science*, 66(9):3843–3859, 2020.
- [116] John Venable, Jan Pries-Heje, and Richard Baskerville. A Comprehensive Framework for Evaluation in Design Science Research. In *International Conference on Design Science Research in Information Systems*, pages 423–438, 2012.
- [117] Friedhelm Victor. Address Clustering Heuristics for Ethereum. In *International Conference on Financial Cryptography and Data Security (FC 2020)*, 2020.
- [118] Friedhelm Victor and Bianca Katharina Lüders. Measuring Ethereum-Based ERC20 Token Networks. In *International Conference on Financial Cryptography and Data Security (FC 2019)*, pages 113–129, 2019.
- [119] Friedhelm Victor, Peter Ruppel, and Axel Küpper. A Taxonomy for Distributed Ledger Analytics. *IEEE Computer Society*, 54, 2021.
- [120] Victor von Wachter, Johannes Rude Jensen, and Omri Ross. How Decentralized is the Governance of Blockchain-based Finance? Empirical Evidence from four Governance Token Distributions. *arXiv*, 2020. URL <https://arxiv.org/abs/2102.10096>.
- [121] Victor Von Wachter, Johannes Jensen, and Omri Ross. Measuring Asset Composability as a Proxy for Ecosystem Integration. In *International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.*, 2021.
- [122] Victor Von Wachter, Johannes Rude Jensen, Ferdinand Regner, and Omri Ross. NFT Wash Trading: Quantifying Suspicious Behaviour in NFT markets. In *International Conference on Financial Cryptography and Data Security (FC 2022). International Workshops.*, 2022.

- [123] Clara Walsh, Philip O'Reilly, Rob Gleasure, Joseph Feller, Shanping Li, and Jerry Cristoforo. New Kid on the Block: A Strategic Archetypes Approach to Understanding the Blockchain. In *International Conference on Information Systems (ICIS 2016)*, 2016.
- [124] Sam M. Werner, Daniel Perez, Lewis Gudgeon, Aariah Klages-Mundt, Dominik Harz, and William J. Knottenbelt. SoK: Decentralized Finance (DeFi). *arXiv*, 2021. URL <https://arxiv.org/abs/2101.08778v6>.
- [125] Douglas West. *Introduction to Graph Theory*. Pearson India, 2015.
- [126] Rolf T. Wigand. Some Recent Developments in Organizational Communication: Network Analysis – A Systemic Representation of Communication Relationships. *comm*, 3(2):181–200, 1977.
- [127] Rolf T. Wigand. Integrated Services Digital Networks: Concepts, Policies, and Emerging Issues. *Journal of Communication*, 38(1):29–49, 1988.
- [128] Maximilian Wohrer and Uwe Zdun. Design Patterns for Smart Contracts in the Ethereum Ecosystem. In *IEEE 2018 International Congress on Cybermatics: 2018 IEEE Conferences on Internet of Things, Green Computing and Communications, Cyber, Physical and Social Computing, Smart Data, Blockchain, Computer and Information Technology*, pages 1513–1520, 2018.
- [129] Gavin Wood. Ethereum: A Secure Decentralized Generalized Transaction Ledger. *Whitepaper*, 2015. URL <https://gavwood.com/paper.pdf>.
- [130] Karl Wüst and Arthur Gervais. Do you need a Blockchain? In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 2018.
- [131] Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) Protocols. *ACM Computing Surveys*, 55(11):1–50, 2023.
- [132] Xiwei Xu, Ingo Weber, Mark Staples, Liming Zhu, Jan Bosch, Len Bass, Cesare Pautasso, and Paul Rimba. A Taxonomy of Blockchain-Based Systems for Architecture Design. In *IEEE International Conference on Software Architecture (ICSA 2017)*, pages 243–252, 2017.
- [133] Dylan Yaga, Peter Mell, Nik Roby, and Karen Scarfone. Blockchain Technology Overview. *National Institute of Standards and Technology*, 2019.
- [134] Dirk A. Zetsche, Ross P. Buckley, Douglas W. Arner, and Linus Ffhr. The ICO Gold Rush: It's a Scam, It's a Bubble, It's a Super Challenge for Regulators. *Harvard International Law Journal*, 60(2), 2019.

- [135] Dirk A Zetzsche, Douglas W Arner, and Ross P Buckley. Decentralized Finance. *Journal of Financial Regulation*, 6:172–203, 2020.
- [136] Rui Zhang, Rui Xue, and Ling Liu. Security and Privacy on Blockchain. *ACM Computing Surveys*, 52(3), 2020.
- [137] Shijie Zhang and Jong Hyouk Lee. Analysis of the Main Consensus Protocols of Blockchain. *ICT Express*, 6(2):93–97, 2020.
- [138] Zibin Zheng, Shaoan Xie, Hongning Dai, Xiangping Chen, and Huaimin Wang. An Overview of Blockchain Technology: Architecture, Consensus, and Future Trends. In *2017 IEEE 6th International Congress on Big Data*, pages 557–564, 2017.
- [139] Liyi Zhou, Xihan Xiong, Jens Ernstberger, Stefanos Chaliasos, Zhipeng Wang, Ye Wang, Kaihua Qin, Roger Wattenhofer, Dawn Song, and Arthur Gervais. SoK: Decentralized Finance (DeFi) Attacks. *arXiv*, 2022. URL <http://arxiv.org/abs/2208.13035>.
- [140] Rafael Ziolkowski, Gianluca Miscione, and Gerhard Schwabe. Decision Problems in Blockchain Governance: Old Wine in New Bottles or Walking in Someone Else’s Shoes? *Journal of Management Information Systems*, 37(2):316–348, 2020.

Paper Overview

The publications are sorted by publication date.

An Introduction to Decentralized Finance (DeFi)

Johannes Rude Jensen, Victor von Wachter, Omri Ross

Abstract Decentralized financial (DeFi) applications are a new breed of consumer-facing financial applications composed of smart contracts, deployed on permissionless blockchain technologies. In this paper, we situate the DeFi concept in the theoretical context of permissionless blockchain technology and provide a taxonomical overview of agents, incentives, and risks. We examine the key market categories and use cases for DeFi applications today and identify four key risk groups for potential stakeholders contemplating the advantages of decentralized financial applications. We contribute novel insights into a rapidly emerging field, with far-reaching implications for financial services.

This paper has been published in *Complex Systems Informatics and Modeling Quarterly Journal* (CSIMQ 2021, 26:3).

2. Measuring Asset Composability as a Proxy for DeFi Integration

Victor von Wachter, Johannes Rude Jensen, Omri Ross

Abstract Decentralized financial (DeFi) applications on the Ethereum blockchain are highly interoperable because they share a single state in a deterministic computational environment. Stakeholders can deposit claims on assets, referred to as *liquidity shares*, across applications producing effects equivalent to rehypothecation in traditional financial systems. We seek to understand the degree to which this practice may contribute to financial integration on Ethereum by examining transactions in *composed* derivatives for the assets DAI, USDC, USDT, ETH, and tokenized BTC for the full set of 344.8 million Ethereum transactions computed in 2020. We identify a salient trend for *composing* assets in multiple sequential generations of derivatives and comment on potential systemic implications for the Ethereum network.

This paper has been published in the Proceedings of the International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.

3. **NFT Wash Trading - Quantifying Suspicious Behavior in NFT Markets**

Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, Omri Ross

Abstract The smart contract-based markets for NFT on the Ethereum blockchain have seen tremendous growth in 2021, with trading volumes peaking at \$3.5b in September 2021. This dramatic surge has led to industry observers questioning the authenticity of on-chain volumes, given the absence of identity requirements and the ease with which agents can control multiple addresses. We examine potentially illicit trading patterns in the NFT markets from January 2018 to mid-November 2021, gathering data from the 52 largest collections by volume. Our findings indicate that within our sample 3.93% of addresses, processing a total of 2.04% of sale transactions triggers suspicions of market abuse. Flagged transactions contaminate nearly all collections and may have inflated the authentic trading volumes by as much as \$149,5m for the period. Most flagged transaction patterns alternate between a few addresses, indicating a predisposition for manual trading. We submit that the results presented here may serve as a viable lower-bound estimate for NFT wash trading on Ethereum. Even so, we argue that wash trading may be less common than what industry observers have previously estimated. We contribute to the emerging discourse on the identification and deterrence of market abuse in the cryptocurrency markets.

This paper has been published in the Proceedings of the International Conference on Financial Cryptography and Data Security (FC 2022). International Workshops.

4. **Blockchain-based Infrastructure for Emerging Economies**

Johannes Rude Jensen, Victor von Wachter, Omri Ross

Abstract The transformative capacity of blockchain technology is a frequently debated topic in the IS and practitioner literature. Yet, rigorous and design-driven research remains relatively uncommon. We document an ongoing design process towards a blockchain-based IT-artefact comprising financial infrastructure for stakeholders in emerging economies. Working with a young NGO, we utilize the DSR in the design, implementation, and evaluation of the IT-artefact. The artifact enables

stakeholders to conduct basic financial services by computing transactions, maintaining a savings account, and receiving targeted stimulus payments. Following six months of iterative design and development, we released a global pilot version of the artifact. Over the first nine months, the pilot generated a dataset of 6.6 million transactions among 189,379 verified users. By conducting design-driven research, we contribute novel practical insights to the IS discourse on the transformative capacity of blockchain technology and information communication technologies (ICT) in emerging economies.

This paper has been published in the Proceedings of the European Conference on Information Systems (ECIS 2022).

5. **Kickstarting Blockchain: Designing Blockchain-based Tokens for Equity Crowdfunding**

Tobias Guggenberger, Benjamin Schellinger, Victor von Wachter, Nils Urbach

Abstract Blockchain-based tokens seek to overcome the friction and opaqueness of the legacy financial infrastructure in the company funding process, particularly in the early-stage and equity crowdfunding domains. While Initial Coin Offerings and Security Token Offerings proposed a solution for crowdfunding, early-stage companies still face challenges in using blockchain as an alternative equity funding infrastructure. In this context, the idea of blockchain-based equity tokens remains hypothetical. In addition, the literature lacks a design theory for the development and implementation of blockchain-based equity tokens. This research bridges this gap by designing, developing, and evaluating an equity token prototype for crowdfunding, following the design science research approach. We propose a refined crowdfunding model and derive seven design principles that contribute to the design theory of equity tokens. The research results show that blockchain-based equity tokens improve efficiency, transparency, and interoperability while meeting regulatory requirements and facilitating secondary market trading.¹

This paper has been published in Electronic Commerce Research Journal (2023).

6. **Fundamentals of Perpetual Futures**

Songrun He, Asaf Manela, Omri Ross, Victor von Wachter

¹For the sake of clarity this paper is a continuation of a Master's thesis. An estimated 66% of the initial work has been revised and produced as part of the Ph.D. process. Significant modifications had been done to various aspects of the initial work, particularly the concept, method, analysis, evaluation, and discussion.

Abstract Perpetual futures, swap contracts that never expire, are the most popular derivative traded in cryptocurrency markets, with more than \$100 billion traded daily. Perpetuals provide investors with leveraged exposure to cryptocurrencies, which does not require rollover or direct cryptocurrency holding. To keep the gap between perpetual futures and spot prices small, long position holders periodically pay short position holders a funding rate proportional to this gap. The funding rate incentivizes trades that tend to narrow the futures-spot gap. But unlike fixed-maturity futures, perpetuals are not guaranteed to converge to the spot price of their underlying asset at any time, and familiar no-arbitrage prices for perpetuals are not available, as the contracts have no expiry date to enforce arbitrage. Here, using a weaker notion of random-maturity arbitrage, we derive no-arbitrage prices for perpetual futures in frictionless markets and no-arbitrage bounds for markets with trading costs. These no-arbitrage prices provide a valuable benchmark for perpetual futures and simultaneously prescribe a strategy to exploit divergence from these fundamental values. Empirically, we find that deviations of crypto perpetual futures from no-arbitrage prices are considerably larger than those documented in traditional currency markets. These deviations comove across cryptocurrencies and diminish over time as crypto markets develop and become more efficient. A simple trading strategy generates large Sharpe ratios even for investors paying the highest trading costs on Binance, which is currently the largest crypto exchange by volume.

This paper is currently under peer review. This version is from April 2023.

An Introduction to Decentralized Finance (DeFi)

Johannes Rude Jensen, University of Copenhagen, eToroX Labs

Victor von Wachter, University of Copenhagen

Omri Ross, University of Copenhagen, eToroX Labs

This paper has been published in Complex Systems Informatics and Modeling Quarterly Journal (CSIMQ 2021, 26:3).

Abstract Decentralized financial applications (DeFi) are a new breed of consumer-facing financial applications composed of smart contracts, deployed on permissionless blockchain technologies. In this paper, we situate the DeFi concept in the theoretical context of permissionless blockchain technology and provide a taxonomical overview of agents, incentives, and risks. We examine the key market categories and use cases for DeFi applications today and identify four key risk groups for potential stakeholders contemplating the advantages of decentralized financial applications. We contribute novel insights into a rapidly emerging field, with far-reaching implications for financial services.

Keywords Blockchain, Decentralized Finance, DeFi, Smart Contracts

Introduction

Decentralized financial applications, colloquially referred to as ‘DeFi’, are a new type of open financial applications deployed on publicly accessible, permissionless blockchains. A rapid surge in the popularity of these applications saw the total value of the assets locked in DeFi applications (TVL) grow from \$675mn at the outset of 2020 to an excess of \$40bn towards the end of the first quarter in the following year¹. While scholars within the information systems

¹<https://defipulse.com/>

and management disciplines recognize the novelty and prospective impact of blockchain technologies, theoretical or empirical work on DeFi remains scarce [13]. In this paper, we provide a conceptual introduction to ‘DeFi’ situated in the theoretical context of permissionless blockchain technology. We introduce a taxonomy of agents, roles, incentives, and risks in DeFi applications and present four potential sources of complexity and risk [11].

Permissionless Blockchain Technology and Decentralized Finance

The implications and design principles for blockchain and distributed ledger technologies have generated a growing body of literature in the information systems (IS) genre [14]. Primarily informed by the commercial implications of smart contract technology, scholars have examined the implications for activities in financial services such as the settlement and clearing of ‘tokenized’ assets [18] the execution and compilation of financial contracts [12, 9, 17], complexities in supply-chain logistics [8] and beyond. A blockchain is a type of distributed database architecture in which a decentralized network of stakeholders maintains a singleton state machine. Transactions in the database represent state transitions disseminated amongst network participants in ‘blocks’ of data. The correct order of the blocks containing the chronological overview of transactions in the database is maintained with the use of cryptographical primitives, by which all stakeholders can manually verify the succession of blocks.

A network consensus protocol defines the rules for what constitutes a legitimate transaction in the distributed database. In most cases, consensus protocols are rigorous game-theoretical mechanisms in which network participants are economically incentivized to promote network security through rewards and penalties for benevolent or malicious behavior [3]. Scholars typically differentiate between ‘permissioned’ and ‘permissionless’ blockchains. Permissionless blockchains are open environments accessible by all, whereas permissioned blockchains are inaccessible to external parties not recognized by a system administrator [14]. Recent implementations of the technology introduce a virtual machine, the state of which is maintained by the nodes supporting the network. The virtual machine is a simple stack-based architecture, in which network participants can execute metered computations denominated in the native currency format. Because all ‘nodes’ running the blockchain ‘client’ software must replicate the computations required for a program to run, computational expenditures are priced on the open market. This design choice is intended to mitigate excessive use of resources leading to network congestion or abuse. Network participants pass instructions to the virtual machine in a higher-level programming language, the most recent generation of which is used to write programs, referred to as smart contracts. Because operations in

the virtual machine are executed in a shared state, smart contracts are both transparent and stateful, meaning that any application deployed as a smart contract executes deterministically. This ensures that once a smart contract is deployed, it will execute exactly as instructed.

DeFi Application Taxonomy

We denote the concept of ‘DeFi application’ as an arrangement of consumer-facing smart contracts, executing a predefined business logic within the transparent and deterministic computational environment afforded by permissionless blockchain technology. Blockchain technology is the core infrastructure layer (see Figure A.1) storing transactions securely and providing game theoretic consensus through the issuance of a native asset. As a basic financial function, standardized smart contracts are utilized to create base assets in the asset layer. These assets are utilized as basis for more complex financial instruments in the application layer. In the application layer, DeFi applications are deployed as sophisticated smart contracts and thus execute a given business logic deterministically. Contemporary DeFi applications provide a range of financial services within trading, lending, derivatives, asset management and insurance services. Aggregators source services from multiple applications, largely to provide the best rates across the ecosystem. Finally, user-friendly frontends combine the applications and build a service similar to today’s banking apps.

In contrast to traditional banking services, in a blockchain-based technology stack, users interact directly with the application independent of any intermediary service provider. The metered pricing of computational resources on permissionless blockchains means that DeFi applications are constrained by the computational resources they can use. Application designers seek to mitigate the need for the most expensive operations, such as storing big amounts of data or conducting sophisticated calculations, in an effort of reducing the level of complexity required to execute the service that their application provides. Because the resources required for interacting with a smart contract are paid by the user, DeFi application designers employ an innovative combination of algorithmic financial engineering and game theory to ensure that all stakeholders of their application are sufficiently compensated and incentivized. In Table A.1, we introduce a taxonomy for the different types of agents and their roles in contemporary DeFi applications. We highlight the incentives for participation and key risks associated with each role. Owing to the original open-source ethos of blockchain technology, application designers are required to be transparent and build ‘open’ and accessible applications, in which users can take ownership and participate in decision-making processes, primarily concerning new features or changes to the applications. As a reaction to these demands, application designers often issue and distribute

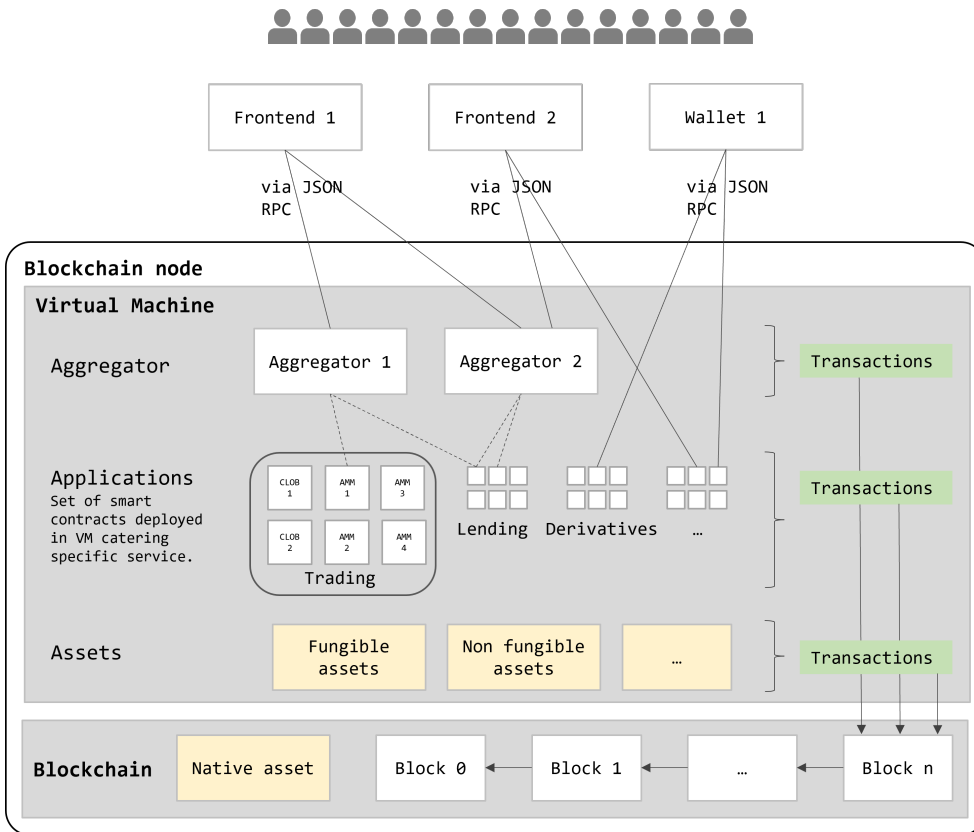


Figure A.1: DeFi applications on permissionless blockchain

so-called governance tokens. Governance tokens are fungible units held by users, which allocate voting power in majority voting schemes [21]. Much like traditional equities, governance tokens trade on secondary markets which introduces the opportunity for capital formation for early stakeholders and designers of successful applications. By distributing governance tokens, application designers seek to disseminate value to community members while retaining enough capital to scale the development of the application by selling inventory over multiple years.

The generalized agent classification demonstrated in Table A.1 is applicable to a wide area of DeFi applications providing peer-to-peer financial services on blockchain technology including trading, lending, derivatives, and asset management. In the following section, we dive into a number of recent use cases, examining the most recently popular categories of applications.

Table A.1 Agent classification, incentives, and key risks

Agent	Role	Incentives for participation	Key risk
Users	Utilizing the application	Profits, credit, exposure, and governance token	Market risk, technical risk
Liquidity Providers	Supply capital to the application in order to ensure liquidity for traders or borrowers	Protocol fees, governance tokens	Systemic economic risk, technical risk, regulatory risk, opportunity costs of capital
Arbitrageurs	Return the application to an equilibrium state through strategic purchasing and selling of assets	Arbitrage profits	Market risk, network congestion, and transaction fees
Application Designers (Team & Founders)	Design, implement, and maintain the application	Governance token appreciation	Software bugs

An Overview of Popular DeFi Application Categories

The development principles presented above have been implemented in a number of live applications to date. In this section, we provide a brief overview of the main categories of DeFi applications.

Decentralized Exchanges and Automated Market Makers Facilitating the decentralized exchange of assets requires an efficient solution for matching counterparties with the desire to sell or purchase a given asset at a certain price, a process known as price discovery. Early implementations of decentralized exchanges on permissionless blockchain technologies successfully demonstrated the feasibility of executing decentralized exchange of assets on permissionless blockchain technology, by imitating the conventional central limit order book (CLOB) design. However, for the reasons stipulated above, this proved infeasible and expensive at scale. First, in the unique cost structure of the blockchain-based virtual machine format [23], traders engaging with an application, pay fees corresponding to the complexity of the computation and the amount of storage required for the operation they wish to compute.

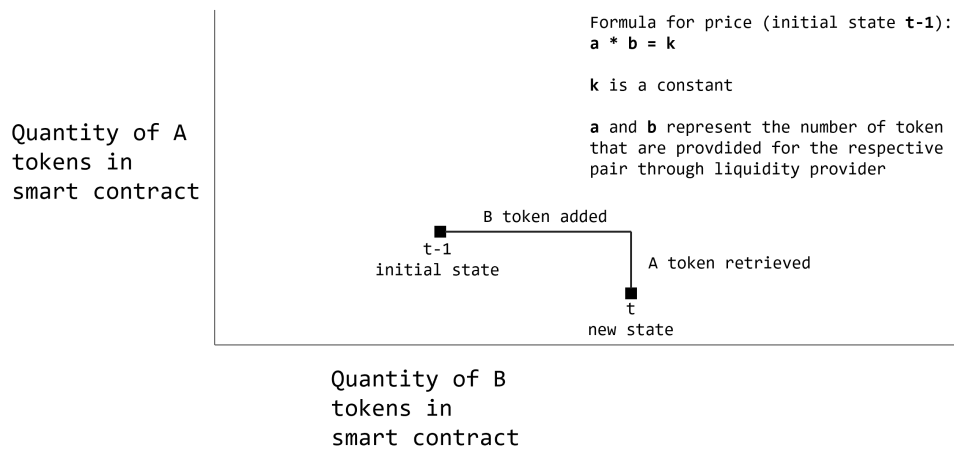


Figure A.2: AMM Price discovery function

Because the virtual machine is replicated on all active nodes, storing even small amounts of data is exceedingly expensive. Combined with the complex matching logic required to maintain a liquid orderbook, computing fees rapidly exceeded users' willingness to trade. Second, as 'miners' pick transactions for inclusion in the next block by the amount of gas attached to the block, it is possible to front-run state changes to the decentralized orderbook by attaching a large computational fee to a transaction including a trade, which pre-emptively exploits the next state change of the orderbook, thus profiting through arbitrage on a deterministic future state [7].

Subsequent iterations of decentralized exchanges addressed these issues by storing the state of the orderbook separately, using the blockchain only to compute the final settlement [22]. Nevertheless, problems with settlement frequency persisted, as these implementations introduced complex coordination problems between orderbook storage providers, presenting additional risk vectors to storage security. Motivated by the shortcomings of the established CLOB design a generation of blockchain-specific 'automated' market makers (AMMs) presents a new approach to blockchain-enabled market design. By pooling available liquidity in trading pairs or groups, AMMs eliminate the need for the presence of buyers and sellers at the same time, facilitating relatively seamless trade execution without compromising the deterministic integrity of the computational environment afforded by the blockchain. Trading liquidity is provided by 'liquidity providers' which lock crypto assets in the pursuit of trading fee returns.

While the primary context for the formal literature on blockchain-based AMM has been provided by Angeris and Chitra et al. [2, 1, 6] the field has attracted new work on adjacent topics such as liquidity provisioning [4, 19, 5]

and token weighted voting systems [20].

Peer-to-peer Lending and Algorithmic Money Markets The ‘money markets’ to borrow and lend capital with corresponding interest payments occupy an important role in traditional financial services. Within DeFi, borrowing and lending applications are amongst the largest segments of financial applications with \$7bn total value locked² at the end of 2020. In borrowing/lending protocols agents with excess capital can lend crypto assets (‘liquidity providers’) to a peer-to-peer protocol receiving continuous interest payments. Consequently, a borrower can borrow crypto assets and pays an interest rate. Given the pseudonymous nature of blockchain technology, it is not possible to borrow funds purely on credit. To borrow funds, the borrowing agent has to ‘overcollateralize’ a loan, by providing other crypto assets exceeding the dollar value of the loan to the smart contract. The smart contract then issues a loan relative to 70-90% of the value of the collateral assets. Should the value of the collateral assets drop below the value of the outstanding loan, the smart contract automatically auctions away the collateral on a decentralized exchange at a profit. The interest rate is algorithmically set by the relative supply and demand for each specific crypto asset. Initially pioneered by the MakerDAO application, several protocols are now accessible providing similar services with novel interest rate calculations or optional insurance properties, currently presiding over a \$7bn crypto assets under management.

Derivatives Blockchain-based financial contracts (derivatives) are one of the fastest-growing market segments in DeFi. Here, application designers seek to make traditional financial derivatives such as options, futures, and other kinds of synthetic contracts available to the broader DeFi ecosystem. A futures contract stipulates a sale of an asset at a specified price with an expiry date, an option contract stipulates the right but not the obligation to sell or purchase assets at a specific price.

As in traditional finance, both financial services can be used as insurance against market movements as well as speculation on prices. Recently, a new branch of ‘synthetic’ assets has entered the market in the form of tokens pegged to an external price, commonly tracking the price of commodities (e.g. gold) or stocks (e.g. Tesla). A user can create such synthetic assets by collateralized crypto assets in a smart contract similar to how decentralized lending is computed. The synthetic asset tracks an external price feed (‘oracle’) which is provided to the blockchain. However, external price feeds are prone to technical issues and coordination problems leading to staleness in case of network congestions or fraudulent manipulation [16].

²<https://defipulse.com/>

Automated Asset Management The traditional practice of ‘asset management’ in the financial services industry consists primarily of the practice of allocating financial assets such as to satisfy the long-term financial objectives of an institution or an individual. As the reader will have noted above, there are an increasing number of DeFi applications, all of which operate algorithmically without human intervention. This means that the DeFi markets operate around the clock and are impossible to manage.

The two main use cases for automated asset managers are ‘yield aggregators’ and traditional crypto asset indices. Utilizing the interoperability and automation of blockchain technology, ‘yield aggregators’ are smart contract protocols allocating crypto-assets according to predefined rules, often with the goal of maximizing yield whilst controlling risk. Users typically allocate assets to a protocol, which automatically allocates assets across applications in order to optimize the aggregate returns, while rebalancing capital allocations on an ongoing basis.

Indices, on the other hand, offer broad exposure to crypto assets akin to the practice of ‘passive investing’. These applications track a portfolio of crypto assets by automatically purchasing these assets and holding them within the smart contract. Equivalent to exchange-traded funds (ETFs), stakeholders purchase ownership of the indices by buying a novel token, granting them algorithmic rights over a fraction of the total assets held within the smart contract.

Identifying and Managing Risk in Decentralized Finance

In this section, we identify and evaluate four risk factors that are likely to introduce new complexities for stakeholders involved with DeFi applications.

Software integrity and security Owing to the deterministic nature of permissionless blockchain technology, applications deployed on smart contracts are subject to excessive security risks, as any signed transaction remains permanent once included in a block. The irreversible or, ‘immutable’ nature of transactions in a blockchain network has led to significant loss of capital on multiple occasions, most frequently as a result of coding errors, sometimes relating to even the most sophisticated aspects of the virtual machine and programming language semantics [15]. DeFi applications rely on the integrity of smart contracts and the underlying blockchain. Risk is further enforced through uncertainties in future developments and the novelty of the technology.

Transaction costs and network congestion To mitigate abusive or excessive use of the computational resources available on the network, computa-

tional resources required to interact with smart contracts are metered. This creates a secondary market for transactions, in which users can outbid each other by attaching transaction fees in an effort of incentivizing miners to select their transaction for inclusion in the next block [7]. In times of network congestion, transactions can remain in a pending state, which ultimately results in market inefficiency and information delays. Furthermore, transaction fees appreciate to an extent to which single applications or sub-components gross several hundreds of thousands of dollars from users interacting with the application³. While intermediary service providers occasionally choose to subsidize protocol transaction fees⁴, application fees are in nearly all cases paid by the user interacting with the DeFi application. Because application designers seek to lower aggregate transaction costs, protocol fees, slippage, or impermanent loss through algorithmic financial modeling and incentive alignment, stakeholders must carefully observe the state of the blockchain network. If a period of network congestion coincides with a period of volatility, the application design may suddenly impose excessive fees or penalties on otherwise standard actions such as withdrawing or adding funds to a lending market [16].

Participation in decentralized governance Responding to implications of the historically concentrated distribution of native assets amongst a small minority of stakeholders, DeFi application designers increasingly rely on a gradual distribution of fungible governance tokens in an attempt to adequately ‘decentralize’ decision-making processes [21]. While the distribution of governance tokens remains fairly concentrated amongst a small group of colluding stakeholders, the gradual distribution of voting power to liquidity providers and users will result in an increasingly long-tailed distribution of governance tokens. Broad distributions of governance tokens may result in adversarial implications of a given set of governance outcomes, for stakeholders who are not sufficiently involved in monitoring the governance process [20].

Application interoperability and systemic risks A key value proposition for DeFi applications is the high level of interoperability between applications. As most applications are deployed on the Ethereum blockchain, users can transact seamlessly between different applications with settlement times rarely exceeding a few minutes. This facilitates rapid capital flows between old and new applications on the network. While interoperability is an attractive feature for any set of financial applications, tightly coupled and complex liquidity systems can generate an excessive degree of financial integration, resulting in systemic dependencies between applications [10]. This factor is exacerbated by the often complex and heterogeneous methodologies for the computation of exposure, debt, value, and collateral value that DeFi

³<https://etherscan.io/gastracker>

⁴<https://coinbase.com>

application designers have used to improve their products. An increasing degree of contagion between applications may introduce systemic risks, as a sudden failure or exploit in one application could ripple throughout the network, affecting stakeholders across the entire ecosystem of applications. The primary example of this dynamic can be demonstrated by the computation of ownership in so-called liquidity pools used by traders utilizing AMM smart contracts. When providing liquidity in the form of crypto assets to a decentralized exchange, liquidity providers receive ‘liquidity shares’ redeemable for a proportional share of the liquidity pool, together with the accumulated fees generated through trading. As liquidity shares are typically transferable and fungible IOU tokens representing fractional ownership of a liquidity pool, this has led to the emergence of secondary markets for liquidity shares. Providing liquidity in the form of IOU tokens, to these secondary markets creates additional (3rd generation) liquidity shares, generating additional fees for the liquidity provider. As a consequence of the increasingly integrated market for liquidity shares, a rapid depreciation of the source asset for the liquidity shares may trigger a sequence of cascading liquidations, as the market struggles to price in any rapid changes in the price of the source asset [16, 10].

Conclusion: Is DeFi The Future of Finance?

In this paper, we have examined the potential implications, complexities, and risks associated with the proliferation of consumer-facing DeFi applications. While DeFi applications deployed on permissionless blockchains present a radical potential for transforming consumer-facing financial services, the risks associated with engaging with these applications remain salient. Future stakeholders contemplating an engagement with these applications ought to consider and evaluate key risks prior to committing or allocating funds to DeFi applications.

Scholars interested in DeFi applications may approach the theme from numerous angles, extending early research on the market design of DeFi applications [1] or issues related to governance tokens [21, 20] and beyond.

Paper References

- [1] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An Analysis of Uniswap Markets. *Cryptoeconomic Systems Journal*, pages 1–25, 2019.
- [2] Guillermo Angeris, Alex Evans, and Tarun Chitra. When does the tail wag the dog? Curvature and Market Making. *arXiv*, 2020. URL <https://arxiv.org/abs/2012.08040>.
- [3] Andreas Antonopoulos and Gavin Wood. *Mastering Ethereum*. O’Reilly Media, Sebastopol, California, USA, 2018.
- [4] Jun Aoyagi. Liquidity Provision by Automated Market Makers Preliminary and Incomplete. *SSRN Electronic Journal*, 2020. URL <https://ssrn.com/abstract=3674178>.
- [5] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Luch-Lafuente. SoK: Lending Pools in Decentralized Finance. *arXiv*, 2020. URL <http://arxiv.org/abs/2012.13230>.
- [6] Tarun Chitra. Competitive Equilibria between Staking and on-chain Lending. *arXiv*, 2019. URL <https://arxiv.org/abs/2001.00919>.
- [7] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. In *IEEE Symposium on Security and Privacy*, pages 910–927, 2020.
- [8] Boris Düdder and Omry Ross. Timber Tracking: Reducing Complexity of Due Diligence by using Blockchain Technology. In *2nd Workshop on Managed Complexity*, 2017.
- [9] Benjamin Egelund-Müller, Martin Elsmann, Fritz Henglein, and Omri Ross. Automated Execution of Financial Contracts on Blockchains. *Business and Information Systems Engineering*, 59(6):457–467, 2017.

- [10] Lewis Gudgeon, Daniel Perez, Dominik Harz, Benjamin Livshits, and Arthur Gervais. The Decentralized Financial Crisis. In *IEEE Crypto Valley Conference on Blockchain Technology*, 2020.
- [11] Johannes Rude Jensen and Omri Ross. Managing Risk in DeFi. In *CEUR Workshop Proceedings*, pages 133–138, 2020.
- [12] Johannes Rude Jensen and Omri Ross. Settlement with Distributed Ledger Technology. In *International Conference on Information Systems (ICIS 2020)*, 2020.
- [13] John Kolb, Moustafa AbdelBaky, Randy H Katz, and David E Culler. Core Concepts, Challenges, and Future Directions in Blockchain. *ACM Computing Surveys*, 53(1):1–39, 2020.
- [14] Olga Labazova. Towards a Framework for Evaluation of Blockchain Implementations. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [15] Loi Luu, Duc Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making Smart Contracts Smarter. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 254–269, 2016.
- [16] Daniel Perez, Sam M Werner, Jiahua Xu, and Benjamin Livshits. Liquidations: DeFi on a Knife-edge. In *International Conference on Financial Cryptography and Data Security (FC 2021)*, 2021.
- [17] Omri Ross and Johannes Rude Jensen. Compact Multiparty Verification of Simple Computations. In *BIR Workshops*, 2018.
- [18] Omri Ross, Johannes Rude Jensen, and Truls Asheim. Assets under Tokenization: Can Blockchain Technology Improve Post-Trade Processing? In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [19] Martin Tassy and David White. Growth Rate of A Liquidity Provider’s Wealth in $XY = c$ Automated Market Makers. *arXiv*, 2020.
- [20] Gerry Tsoukalas and Brett Hemenway Falk. Token-Weighted Crowdsourcing. *Management Science*, 66(9):3843–3859, 2020.
- [21] Victor von Wachter, Johannes Rude Jensen, and Omri Ross. How Decentralized is the Governance of Blockchain-based Finance? Empirical Evidence from four Governance Token Distributions. *arXiv*, 2020. URL <https://arxiv.org/abs/2102.10096>.

- [22] W. Warren and A. Bandeau. 0x: An open Protocol for Decentralized Exchange on the Ethereum Blockchain, 2017. URL https://github.com/OxProject/whitepaper/blob/master/0x_white_paper.pdf.
- [23] Gavin Wood. Ethereum: A Secure Decentralized Generalized Transaction Ledger, 2015. URL <https://gavwood.com/paper.pdf>.

Measuring Asset Composability as a Proxy for DeFi Integration

Victor von Wachter, University of Copenhagen
Johannes Rude Jensen, University of Copenhagen, eToroX Labs
Omri Ross, University of Copenhagen, eToroX Labs

This paper has been published in the Proceedings of the Financial Cryptography and Data Security Conference (FC 2021). International Workshops.

Abstract Decentralized financial (DeFi) applications on the Ethereum blockchain are highly interoperable because they share a single state in a deterministic computational environment. Stakeholders can deposit claims on assets, referred to as *liquidity shares*, across applications producing effects equivalent to rehypothecation in traditional financial systems. We seek to understand the degree to which this practice may contribute to financial integration on Ethereum by examining transactions in *composed* derivatives for the assets DAI, USDC, USDT, ETH, and tokenized BTC for the full set of 344.8 million Ethereum transactions computed in 2020. We identify a salient trend for *composing* assets in multiple sequential generations of derivatives and comment on potential systemic implications for the Ethereum network.

Keywords Blockchain, DeFi, Asset Composability, Ethereum

Introduction

Smart contracts on the Ethereum blockchain share a single state in a deterministic execution environment [3], a feature that introduces a high level of interoperability between decentralized financial (DeFi) applications. This novelty has thus far, resulted in a rich ecosystem of financial applications, primarily led by borrowing/lending money markets [4, 8] and constant function market makers (CFMM) [1, 2]. At the time of writing, crypto assets valued

in excess of \$39 billion are managed by some 75¹ DeFi applications on the Ethereum blockchain.

From the consumers’ perspective, interoperability between financial applications is a desirable feature, resulting in a vibrant and highly competitive marketplace of increasingly exotic financial products. Yet, if left unsupervised, interoperability between liquidity reserves may lead to dependencies amongst applications, as techniques equivalent to the practice of rehypothecation in the traditional financial system [13] become normalized.

When allocating assets to a CFMM such as Uniswap, Curve, or Balancer, liquidity providers receive ‘liquidity provider shares’ (LP shares) [6] redeemable for a proportional share of the liquidity pool with the unrealized returns of the position. LP shares are typically computed as transferable, fungible tokens which has led to the emergence of new secondary markets in which applications offer liquidity and lending pools for LP shares themselves. Supplying LP shares to these pools results in the issuance of meta LP shares. This process is, in some cases, repeated recursively as stakeholders seek to maximize yield or functionality across a diverse set of applications. While LP shares are often treated by market participants as simple IOUs, they do in fact represent a complex payout function, as shown in the literature by [6, 7]. Further complicating matters, the practice of ‘yield farming’, i.e. allocating assets across DeFi applications to maximize returns [2], has introduced a competitive environment in which applications seek to attract additional liquidity by rewarding LP shareholders with ‘governance tokens’ [12].

We approach Ethereum as a financial ecosystem with structural properties comparable to those of a single market [5, 10]. For this work, we examine the degree to which a crypto asset can be utilized in a sequence of increasingly complex ‘wrapping’ operations, guiding our research question: *Can we measure assets composability as a proxy for financial integration on the Ethereum Blockchain?* Informed by the process proposed by [9], we measure the degree to which crypto assets in smart contracts may contribute towards effects equivalent to financial integration on the Ethereum blockchain. We approach transaction data on Ethereum with an asset-oriented perspective, in contrast to previous studies of financial activity on Ethereum, sorting by addresses [9] or applications [11].

Method

We measure asset composability by identifying the number of derivatives produced from an initial root asset I . We extend the work presented in [9] by proposing an algorithm for unwrapping crypto assets. The algorithm builds a tree structure of derivatives from the initial asset I (Figure B.1). We measure the distance δ to the initial asset $\delta_A = \sum_{i=0}^N |w_i|$ as a proxy for the degree to

¹defipulse.com, as of 31st Jan 2020

```

1: repeat
2:    $T \leftarrow$  all transactions of initial assets from block #9193266 to #11565018
3:   draw 10,000 random transactions  $t$  in  $T$ 
4:   for each  $t$ :
5:     identify ERC20 tokens in transaction
6:     if token  $A$  is wrapped version of initial asset
7:       if  $A$  less than 100 transfers ignore
8:       else  $w_i + 1$  and calculate distance  $\delta_A$ 
9:     end if
10:  end for
11: until no relevant new wrapped assets

```

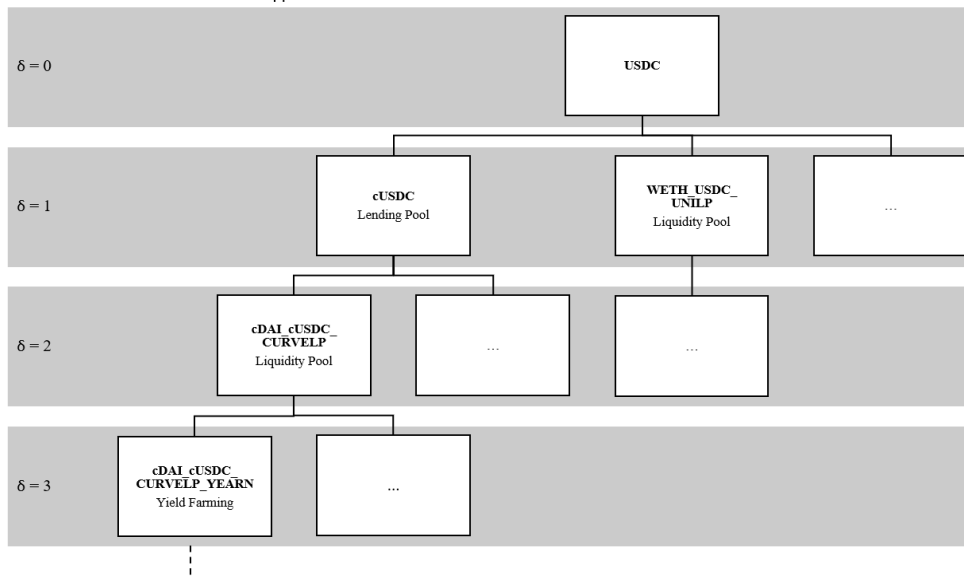


Figure B.1: Method and exemplary asset tree structure for USDC

which an asset contributes towards integration on Ethereum. That is the sum of relevant wrapping operations, where $w := (w_1, \dots, w_n)$ is the vector of all adjustments for the composed asset A .

In the example (Figure B.1), an asset is allocated to a CFMM liquidity pool, triggering the issuance of the corresponding LP shares. At this point, we consider the initial asset as wrapped once, resulting in a distance of 1. Subsequently allocating the LP share to another application would trigger the issuance of another LP share, which amounts to a distance of $\delta = 2$. We target five popular crypto assets: DAI, USDT, USDC, ETH, and tokenized BTC² for the duration of 2020 (Table B.1). Collectively, the selected assets amounted to over 70% of the total value administered within DeFi applications³ at the end of the sample period.

²Bitcoin (BTC) is a non-native asset on Ethereum, represented by ‘wrapped bitcoin’ locked on the original blockchain. We compile the three largest representations of Bitcoin on Ethereum into a single category, assigning the category an initial distance of one.

³defipulse.com, as of 31st Jan 2020

Table B.1 Transactions of plain asset and composed versions during 2020

Initial asset	Transactions on Ethereum	Transactions of composed assets
DAI	4,149,654	1,033,674
USDT	64,956,383	687,705
USDC	7,053,402	1,167,163
WETH	21,187,823	919,165
BTC (wBTC, renBTC, sBTC)	658,035	193,394

Results

We find derivatives of the five initial assets among all 344.8 million Ethereum transactions in 2020 (block #9193266 to #11565018). For each initial asset, we compare the number of transactions in the ‘plain’ version of the asset, against the number of transactions in its derivatives (Figure B.2).

For the first 6 months plain DAI transfers amounted between 82% - 91% (blue) of all DAI asset transfers and composed DAI with $\delta = 1$ amounted between 9% - 18% (orange) respectively. The data indicate a clear trend towards increasingly complex wrapping operations peaking in the third quarter of 2020, a period colloquially referred to as ‘DeFi Summer’ due to a high volume of governance tokens issued at the time [12]. The tendency is especially salient in ‘DAI’, for which up to 84% of all transactions involved a ‘wrapped’ derivative of the initial asset. Curiously, the asset with the largest market cap on Ethereum, USDT, appears to be the least popular with an insignificant 687,705 transactions in ‘wrapped’ derivatives, compared to 64,956,383 transactions in the plain asset.

Discussion and Conclusion

Computing fractional ownership claims in a deterministic, single-state environment introduces a large set of new opportunities for innovation in the financial sector. Because transactions on permissionless blockchains, such as Ethereum, settle atomically, the role of central clearing counterparties in mitigating counterparty risk is largely mitigated for simple transactions. Yet, to date, little is understood about the systemic implications of the design of these applications and how novel concepts like LP shares, may exacerbate the impact of shocks triggered by exploits^{4 5}.

⁴<https://cointelegraph.com/news/akropolis-defi-protocol-paused-as-hackers-get-away-with-2m-in-dai>, accessed 20th Dec 2020

⁵<https://defirate.com/imbtc-uniswap-hack/>, accessed 20th Dec 2020

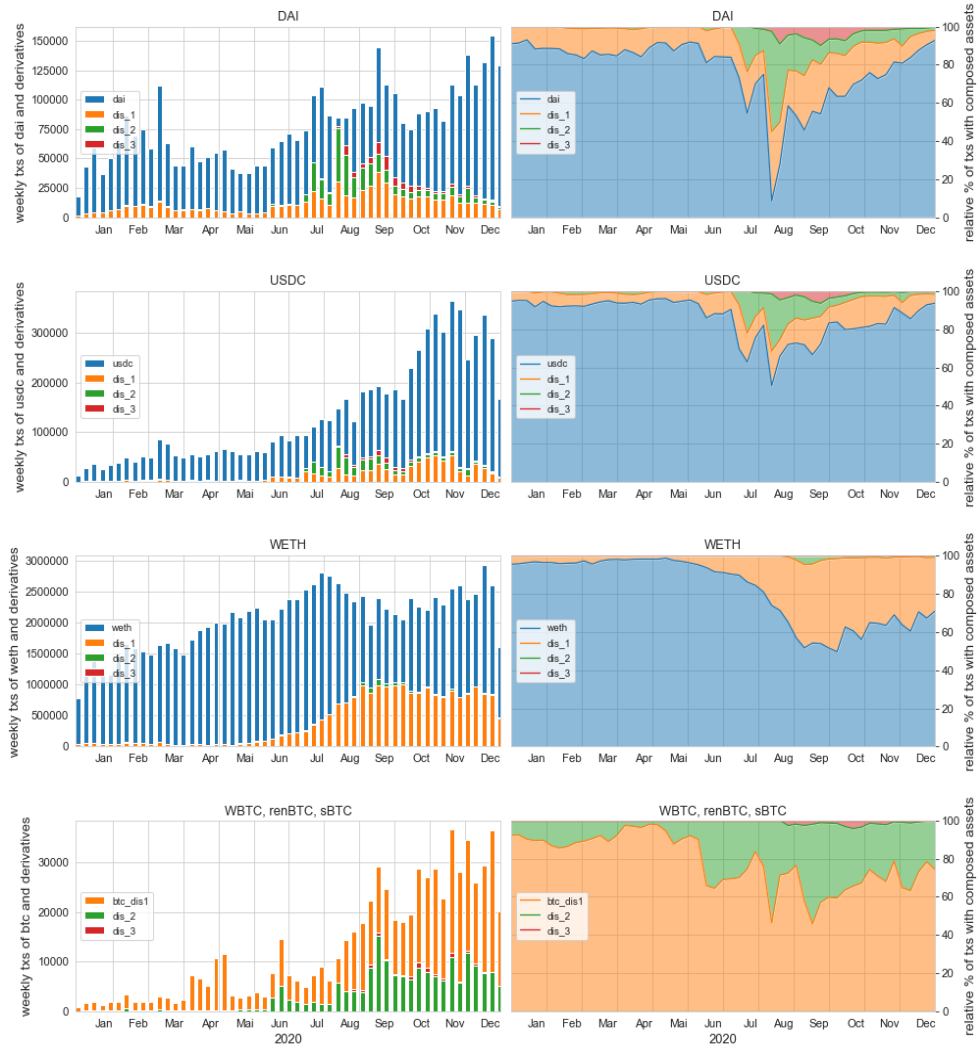


Figure B.2: Financial integration of assets during 2020

A quantifiable approach to the study of financial integration on the Ethereum network will facilitate a better understanding of how shocks travel through tightly interconnected webs of DeFi applications, which may provide guidance toward promoting resilience and protecting investors against systemic risk. In this work, we present initial indicators by examining the degree to which transactions in ‘wrapped’ derivatives of an asset, representing increasingly complex payout functions, may offer an indication of the degree of financial integration on the network. We position this contribution within the broader literature on the quantification of ‘composability risk’ for the DeFi ecosystem, a critical gap raised by [13].

To provide actionable insights for market participants and regulators, this and future studies must expand the scope by considering all relevant factors for the transmission of shocks, including smart-contract design and default risk for the individual DeFi application.

Paper References

- [1] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An Analysis of Uniswap Markets. *Cryptoeconomic Systems Journal*, pages 1–25, 2019.
- [2] Guillermo Angeris, Alex Evans, and Tarun Chitra. When does the Tail wag the Dog? Curvature and Market Making. *arXiv*, 2020. URL <https://arxiv.org/abs/2012.08040>.
- [3] Andreas Antonopoulos and Gavin Wood. *Mastering Ethereum*. O’Reilly Media, Sebastopol, California, USA, 2018.
- [4] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch Lafuente. SoK: Lending Pools in Decentralized Finance. In *International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.*, pages 553–578, 2021.
- [5] Fabio Castiglionesi, Fabio Feriozzi, and Guido Lorenzoni. Financial Integration and Liquidity Crises. *Management Science*, 65(3):955–975, 2017.
- [6] Alex Evans. Liquidity Provider Returns in Geometric Mean Markets. *arXiv*, 2020. URL <https://arxiv.org/abs/2006.08806>.
- [7] Johannes Rude Jensen, Mohsen Pourpouneh, Kurt Nielsen, and Omri Ross. The Homogenous Properties of Automated Market Makers. *arXiv*, 2021. URL <https://arxiv.org/abs/2105.02782>.
- [8] Hsien-Tang Kao, Tarun Chitra, Rei Chiang, and John Morrow Gauntlet. An Analysis of the Market Risk to Participants in the Compound Protocol, 2019. URL https://scfab.github.io/2020/FAB2020_p5.pdf.
- [9] Matthias Nadler and Fabian Schär. Decentralized Finance, Centralized Ownership? An Iterative Mapping Process to Measure Protocol Token Distribution. *arXiv*, 2020. URL <http://arxiv.org/abs/2012.09306>.
- [10] Shahar Somin, Yaniv Altshuler, Goren Gordon, Alex Sandy ‘Pentland’, and Erez Shmueli. Network Dynamics of a Financial Ecosystem. *Scientific Reports*, 10(1), 2020.

- [11] Palina Tolmach, Yi Li, Shang Wei Lin, and Yang Liu. Formal Analysis of Composable DeFi Protocols. In *International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.*, pages 149–161, 2021.
- [12] Victor von Wachter, Johannes Rude Jensen, and Omri Ross. How Decentralized is the Governance of Blockchain-based Finance? Empirical Evidence from four Governance Token Distributions. *arXiv*, 2020. URL <https://arxiv.org/abs/2102.10096>.
- [13] Sam M. Werner, Daniel Perez, Lewis Gudgeon, Aariah Klages-Mundt, Dominik Harz, and William J. Knottenbelt. SoK: Decentralized Finance (DeFi). *arXiv*, 2021. URL <https://arxiv.org/abs/2101.08778v6>.

NFT Wash Trading - Quantifying Suspicious Behavior in NFT Markets

Victor von Wachter, University of Copenhagen
Johannes Rude Jensen, University of Copenhagen, eToroX Labs
Ferdinand Regner, University of Vienna
Omri Ross, University of Copenhagen, eToroX Labs

This paper has been published in the Proceedings of the Financial Cryptography and Data Security Conference (FC 2022). International Workshops.

Abstract The smart contract-based markets for nonfungible tokens (NFTs) on the Ethereum blockchain have seen tremendous growth in 2021, with trading volumes peaking at \$3.5b in September 2021. This dramatic surge has led to industry observers questioning the authenticity of on-chain volumes, given the absence of identity requirements and the ease with which agents can control multiple addresses. We examine potentially illicit trading patterns in the NFT markets from January 2018 to mid-November 2021, gathering data from the 52 largest collections by volume. Our findings indicate that within our sample 3.93% of addresses, processing a total of 2.04% of sale transactions triggers suspicions of market abuse. Flagged transactions contaminate nearly all collections and may have inflated the authentic trading volumes by as much as \$149,5m for the period. Most flagged transaction patterns alternate between a few addresses, indicating a predisposition for manual trading. We submit that the results presented here may serve as a viable lower-bound estimate for NFT wash trading on Ethereum. Even so, we argue that wash trading may be less common than what industry observers have previously estimated. We contribute to the emerging discourse on the identification and deterrence of market abuse in the cryptocurrency markets.

Keywords DeFi, NFT, Blockchain, Wash trading, Graph analysis

Introduction

A nonfungible token (NFT) is a unique digital representation of a digital or physical asset. While the NFT standard is used widely to designate ownership of artifacts such as domain name registrations or concentrated liquidity positions in constant function market makers (CFMM) [22], the arguably most recognized use of the NFT standard is within the representation and trade of digital art and collectibles. Here, the NFT is typically used to represent the ownership of a digital image externally stored, either on a server or, more commonly, on censorship-resistant distributed file systems such as the Interplanetary File System. A basic NFT standard such as the ERC721 [11] typically denotes an interface implementing the ability to own, transfer and trade the NFT. Standardization led to the emergence of NFT markets, facilitating primary and secondary trading, the presently most dominant of which is OpenSea. Permissionless NFT markets, themselves implemented as smart contracts, enable users to sell and purchase NFTs in two ways: as a fixed-price sale or auction, in which competing bids are locked by the smart contract together with the NFT until a winner is found and the auction is cleared. With the admission of NFTs into popular culture, trade volumes on these markets have seen dramatic growth from a mere \$12m settled in September 2020, to volumes exceeding \$3.5b in September of the following year, a surge of over 29,060%¹.

Users typically connect to the permissionless markets through public-key cryptography capable of generating an arbitrary number of addresses [15]. As a consequence, user identities remain entirely pseudonymous in NFT markets, making the obfuscation of illicit practices challenging to prevent. As the unique properties of the Ethereum blockchain simplify adversarial agents to hide in plain sight, we hypothesize that wash trading and strategic bidding amongst multiple addresses controlled by a single or colluding agent may be a frequent occurrence. As there are no theoretical limits to the number of pseudonymous addresses a single agent can control, we conjecture that adversarial agents likely employ a mixture of manual trading and bots to trade NFTs between clusters of addresses in their control. This behavior serves the strategic purpose of artificially inflating the trade volume of a given NFT, creating an impression of desirability to uninformed traders [14]. The uninformed traders, looking for a great opportunity to buy a ‘hot’ NFT will interpret the transaction volume as an authentic expression of interest from other collectors and immediately place a bid or purchase the NFT at an artificially inflated price. Furthermore, novel markets tend to be driven by a volatile search for suitable pricing models [17]. This is undoubtedly the case for the blossoming crypto markets on which the current level of ‘irrational exuberance’ may result in inefficient markets given the presence of uninformed traders looking to strike gold. Adversarial market participants have been shown to exploit

¹<https://dune.xyz/sophieqgu/NFT-Marketplaces>, accessed Dec 2022

these conditions, primarily by employing strategic wash trading on centralized and decentralized central limit orderbook (CLOB) exchanges [1, 8, 19, 24].

Yet, the extent to which these or equivalent practices are being used on NFT markets, remains unclear. To fill the gap, in this paper, we study activities between addresses participating in the NFT markets on Ethereum. We pursue the research question: ‘To what extent does wash trading occur in smart contract-based NFT markets on Ethereum, and to which extent does this practice distort prices?’. Conceptualizing trading patterns as a graph and proposing two detection algorithms, we identify 2.04% as the lower bound of suspicious sale transactions that closely follow the general definition of wash trading.

Literature Review

Wash trading is a well-known phenomenon in traditional financial markets and refers to the activity of repeatedly trading assets for the purpose of feeding misleading information to the market [4]. Typically, one or more colluding agents conduct a set of trades, without taking market risks, which leads to no change in the initial position of the adversarial agents. Most of the early academic publications on wash trading in financial markets focused on colluding investor behavior (e.g. [12]). Cao et al. [4] were among the first to analyze wash trading by specifying trading patterns. Later, they extended their study using directed graphs on order book data [5]. The literature on the identification of wash trading patterns in the cryptocurrency markets primarily emphasizes CLOB models on decentralized exchanges [24] and centralized exchanges, where wash trading practices have been shown to be especially prevalent [1, 8, 19]. Perhaps because the introduction of CFMMs has nearly eliminated the efficacy of wash trading in smart contract-based markets for fungible assets, research on NFTs tends to emphasize either market dynamics and pricing [18, 10] or the technical design considerations [20, 26]. Thus far, little academic research has examined market abuse in smart contract-based NFT markets [9].

Methodology

The Ethereum blockchain is a type of permissionless ledger, in which all transactions and state changes introduced by smart contracts are replicated across all participating nodes in the network [2]. This introduces a high level of integrity to the database, but simultaneously requires pseudonymization, as anyone with access to the database would otherwise be able to view the balances of users on the network. In Ethereum, this problem is solved with public-key cryptography [2]. Any user on the network can generate a public/private key pair, which can subsequently be used to generate an arbitrary

number of addresses. This design presents a fascinating paradox from the perspective of identifying market abuse: Pseudonymous identities are essential in protecting the privacy of benevolent users but, at the same time, they allow adversarial agents to hide in plain sight. Yet, due to the strict ordering of transactions and unique properties of NFT markets, blockchain transaction data presents a powerful and unique opportunity for pattern detection [25] utilizing graph-based algorithms [24, 6, 27] and address clustering [23, 13].

Data Aggregation and Cleaning: We collect transaction data on the 52 leading ERC721 NFT collections on the Ethereum blockchain by trading volume, covering a period between the 1st of January 2018 until 21st of November 2021. The dataset contains 21,310,982 transactions of 3,572,483 NFTs conducted by 459,954 addresses. Collectively, the dataset represents \$6.9b of the \$12.3b total trading volume (49.5%)² on all NFT markets since the first block of the Ethereum blockchain. We capture all blockchain transactions related to the selected NFT collections via the OpenSea API. We parse the dataset by ‘sale’ events emitted by an NFT contract when it is transacted on a smart contract-based marketplace, indicating that a change of ownership has been recorded on the blockchain, and ‘transfer’ events indicating that the NFT has been transferred from one address to another. The dataset was subsequently enriched with (I) historic USD prices for settlements in crypto and stablecoins via the Coingecko API and (II) blockchain-specific data via the Etherscan services. (III) To maintain an accurate overview of the NFT markets in the dataset, we collected the deployment date of the four largest NFT on-chain markets manually (Foundation, OpenSea, Rarible, and Superrare) and matched the deployment dates with the event emissions. It should be noted that the dataset collected for this analysis pertains only to operations conducted within the Ethereum Virtual Machine (EVM), meaning that we knowingly omit any ‘off-chain’ transactions or bidding patterns from the analysis. Finally, we pre-processed the dataset with standardized scripts, eliminating a very small fraction of transactions due to obvious technical errors or trades against exotic assets for which the price data tends to be inaccurate.

Building Transaction Graphs: In order to identify suspicious behavior conforming with wash trading activity, we model the transaction history of each NFT as a directed multigraph $G_{nft} = (N, E)$, where N is the set of addresses and E is the set of ERC721 transactions between addresses. The direction of the edges is given by the transaction flow from sender to receiver, identified by the transaction hash. The weight of the edges represents the USD price at the time of the transaction. This denotation is amenable to the identification of clusters in which a sequence of transactions leads to no apparent position change for any of the addresses involved. Topologically, these patterns form closed cycles. Utilizing Deep-First-Search-Algorithm [21] we identify closed cycles within the data set. We adjust and iterate the

²<https://dune.xyz/sophieqgu/NFT-Marketplaces>, accessed Dec 2022




Figure C.1: Detection of suspicious activity through closed cycles. We exclude cycles involving only ‘transfer’ transactions without taking market events.

Figure C.2: Detection of suspicious activity through a rapid sequence of transactions without taking market risk.

algorithm, using the temporal distance between transactions to detect sub-cycles. The proposed algorithm (Appendix) has a linear time complexity of $\mathcal{O}((n + e)(c + 1))$ for n nodes (addresses), e edges (transactions) and c cycles [16]. Figure C.1 illustrates a few examples of suspicious cyclic activities. Transactions E belonging to a suspicious cycle are marked in red, and potentially colluding trading addresses N , are highlighted in grey. Example 1 is a self-directed transaction. Examples 2-4 illustrate variations of a cycle with two transactions, where solid lines represent ‘sales’ and dotted lines represent ‘transfers’. Example six depicts a complex graph where edges $\{B, C\}$ and $\{B, C, D\}$ form two sub-cycles distinct by time. Further, we analyze path-like transaction patterns, as agents could actively avoid closed cycles. Informed by the definition of wash trading, we consider rapid trade sequences without exposure to market risk as potentially suspicious. Erring on the side of caution, we apply relatively strict thresholds. First, we define the transaction velocity for a sequence as the time elapsed from the initiation to the end. We delimit a rapid trade sequence below 12 hours. Second, we delimit the deviation in USD values, a proxy for market risk, in a sequence to a maximum of 5% of the initial price. Combining both threshold flags 0.3% of the sale transactions as mildly suspicious. Figure C.2 showcases these path-like trade sets.

The proposed algorithms are highly applicable, in that NFT marketplaces deviate from conventional markets in several ways: First, as NFTs are uniquely identifiable by smart contract address and id, detection does not require volume matching required for fungible tokens (e.g. [5, 24]). Second, in contrast to other market designs such as CLOBs, the seller can retain certain control over the opposing counterparty, making it potentially easier to conduct cyclical trades. Lastly, due to the transparency of the Ethereum blockchain, we can inspect trading behavior at the account level without relying on statistical indicators [19].

Table C.1 Overview of the results.

	Dataset	Identified	Percentage
Addresses	459,954	18,117	3.93%
Transactions	1,779,380	36,385 (cyclic: 30,467 sequential: 5,918)	2.04%
Volume in \$	6.9 b	149.5 m	2.17%
NFTs	3,572,483	16,289	0.45%

Results

The analysis flags a total of 3.93% of the addresses as suspicious, indicating that these addresses might be controlled by single agents and used to conduct cyclical or sequential wash trading with NFTs. The flagged addresses processed 2.04% of the total sale transactions, inflating the trading volume by \$149.5m or 2.17% for the period. Of the 36,385 flagged sale transactions, 30,467 were conducted in clusters of cyclical patterns whereas 5,918 were conducted as a rapid sequence. The suspicious activity was executed with just 0.45% of the NFTs in the dataset, indicating a high concentration of illicit activities around a few NFTs (Table C.1). While we identify suspicious activities in all NFT collections (Table C.2, Appendix), the extent to which a collection is contaminated by flagged transactions ranges from 0.19% to 60.93%, indicating that adversarial agents tend to target specific collections for illicit practices. In general, we observe a predisposition for simple trading patterns. 60.6% of the identified clusters are simple variations with two transactions (equivalent to examples 2-4 in Figure C.1). Complex variations of three (8.7%) or more than three transactions are less common (30.7%). However, we find no signs of self-directed trades. Cyclical patterns are conducted at relatively rapid intervals. Figure C.3 illustrates the elapsed time from the first to the last transactions, with respect to the number of transactions involved. Overall, 48.1% of the identified cycles happen within a single day. 13.2% happen within one to seven days and 13.0% are just below 30 days. Consequently, 74.3% are conducted within 30 days, an important threshold under US regulation³. The identified 2-transactions variations have a median execution time of 4.2h (3-transactions: 54h), suggesting a preference for simple and fast patterns. We assume this to indicate that adversarial agents are not trading in an automated fashion, which would result in more complex patterns and execution times within a few minutes. Rapid sequential trades, in which a NFT is moved fast between accounts without any market risk, contributed to only 5918 suspicious transactions, equivalent to 0.3% of the transactions across all collections.

³<https://www.cftc.gov/LawRegulation/CommodityExchangeAct/index.htm>




Figure C.3: Elapsed time to close a cycle with respect to the number of transactions involved. 48.1% of the identified cycles happen within a single day.

Further, we analyze, at what point within the collections’ lifetime suspicious behavior is most prominent. For each collection, we determine the mean suspicious activity starting with the creation date of the respective collection. Figure C.4 (Appendix) shows a peak of suspicious activity in the first third of a collection’s lifetime, possibly in order to raise initial awareness to attract naive buyers. In absolute terms, wash trading is the highest at the beginning of a collection’s lifetime, however, this is also matched by a high amount of organic traffic. Increasing the price of an asset by faking activity is a central motivation for agents in conducting illicit trading [14]. Analyzing if the average price is inflated through wash trading practices, we find that the subsequent sale after a detected wash trade has, on average, an increased price of 30.53%. However, a regression on panel data to measure the impact on the price led to insignificant results for a majority of collections. Looking into external factors, we found that *age* has a strong positive relationship with the price. While we expected an inverse relationship of the gas price, which impacts the costs per transaction, with the NFT price, we found the effect to be mixed in a majority of cases. We suspect these findings are influenced by the strong bull market, given the current overall positive public sentiment.

Finally, we explore the relationship between executed trades per address and unique trade partners per address. On NFT markets sellers can retain certain control over the opposing counterparty, thus a large number of trades with only a few other addresses raises suspicion. Figure C.5 (Appendix) visualizes this relationship, whereas addresses that conducted many trades with only a few other addresses would be tilted to the left. Each dot represents an address trading on NFT markets, positioned by the number of trades and

unique trade partners. The size of the dot depicts the number of empirically identified suspicious trades. We find a cluster of suspicious addresses, conducting 25-37 trades with only 12-17 unique trade partners. Furthermore, in contrast to other markets, addresses have relatively few trades, again suggesting a low level of automated trading in this nascent market.

Discussion

Given the challenges in interpreting pseudonymous blockchain-based data, this study has multiple limitations. First and foremost, it should be made clear that none of the findings presented in this paper present any conclusive evidence of criminal activities or malicious intent. While we delimit a set of behaviors that we find unlikely to be conducted with benevolent intent, we leave it to the reader to assess the likelihood that flagged transactions constitute attempts at wash trading. Any or all of the flagged sequences may be erroneous but authentic transactions. Second, the decision of limiting our analysis to a specific subset of cyclical and sequential patterns may result in false negatives as sophisticated attempts at wash trading involving advanced address clusters over longer periods are not flagged by the analysis, at this point. Wash traders may be more careful and evade the analyzed heuristics. Similarly, the analysis pertains exclusively to on-chain transaction events emitted by the NFT contract and does not account for strategic bidding practices, which we suspect may be a popular methodology amongst adversarial agents. Because of these limitations, we hypothesize that the results presented here detect a lower bound for the actual extent of adversarial behavior on decentralized NFT markets.

Should wash trading conducted according to the patterns explored in this paper increase over time, smart contract-based NFT platforms may consider the implementation of obligatory or voluntary identification initiatives. Alternatively, trading limitations on trading velocity, price deltas, or counterparties can be implemented. In our sample self-directed trades have been non-existent, with indicates successful countermeasures at the smart contract or frontend level. Nevertheless, any such attempt at introducing restrictions or limitations may stifle organic market activity and will inevitably create a cat-and-mouse game, as developers and wash traders race to identify and create increasingly sophisticated patterns. More subtle countermeasures fostering the supervision of NFT markets are the expansion of NFT standards beyond ERC721 and ERC1155, as well as increased data ubiquity. Decentralized NFT markets are transparent, however, NFT data is very diverse and difficult to retrieve. Fees potentially play a big role in preventing wash trading, as long as the rewards or incentives are less than the cost of attack. Fees on NFT markets are substantial. Fraudulent agents are less likely to perform a wash trade if they are losing several percentages with every transaction.

Admittedly, this does neither stop marketplaces itself to perform wash trades nor prevents private offset agreements between the trader and a marketplace.

Even so, with \$149.5m and a median of 2.04% suspicious sale transactions we argue that wash trading may be less common than what industry observers have previously estimated [7, 3].

Conclusion

We identify what we believe may serve as a lower bound estimation for suspicious trading behavior on decentralized NFT markets, following the definition of wash trading: sets of trades between collusive addresses, without taking market risk, that lead to no change in the individual position of the participating addresses. Our findings indicate that (I) adversarial agents exhibit a clear preference towards fast and simple cyclical patterns, (II) the level of suspicious activity varies significantly across NFT collections, (III) illicit activity could still be done in a manual fashion, and (IV) the activities do not necessarily produce the intended price impact, as other exogenous factors such as *age* and *sentiment* are more relevant to price discovery.

As a theoretical contribution, we add descriptive knowledge to an emerging field of research where scientific studies are scarce. We contribute to the growing literature on the identification of illicit market behavior in centralized and decentralized crypto markets, by conducting the first in-depth examination of NFT wash trading on the Ethereum blockchain. We contribute empirical statistics of fraudulent behavior and a set of suspicious transaction graphs to foster the understanding of wash trading in increasingly financialized NFT markets. The valuable insights we generate for practitioners are twofold: First, we provide valuable insights to prevent collectors from buying NFTs that are potentially inflated by wash trading. Second, we discuss practical countermeasures to increasing the standards for the wider NFT ecosystem. Further research opportunities are manifold and include studying NFT markets incentivizing volume through tokens, the utilization of flash loans for wash trading, as well as researching the correlation between suspicious behavior and sentiment data.




Figure C.4: Wash trading with respect to collections' lifetime. In absolute terms, wash trading is the highest at the beginning of a collection's lifetime, however, this is also matched by a high amount of organic traffic.

Appendix

Algorithm

Algorithm 1 The detection algorithm

- 1: Input: T timestamped blockchain transactions
 - 2: $L \leftarrow$ empty list of *cycles*
 - 3: **for** $n_{ft} \in T$ **do**
 - 4: $G_{n_{ft}} \leftarrow (N, E)$
 - 5: $G_{n_{ft}} \leftarrow$ identifier, weight
 - 6: label $n \in N$ as discovered
 - 7: **for** all directed E of n **do**
 - 8: test for adjacent edges m
 - 9: **if** m is not labeled as discovered **then**
 - 10: continue
 - 11: **else**
 - 12: $L \leftarrow$ *cycle*
 - 13: $G_{n_{ft}^*} \leftarrow G_{n_{ft}} - E$
 - 14: break and recurse
 - 15: **end if**
 - 16: **end for**
 - 17: **end for**
 - 18: return L
-

Table C.2 Results for each collection. Column (A) is the share of suspicious addresses, (B) the share of suspicious transactions, (C1) represents the total volume flagged denominated in USD, (C2) the share of the flagged volume and (D) the share of suspicious NFTs

	Collection	Smart contract	Created	Size	(A)	(B)	(C1)	(C2)	(D)
1	0n1-force	0x3bf2...5e9d	08/2021	7776	4.2	1.6	1732469	1.2	1.7
2	acclimatedmooncats	0xc3f7...5e69	04/2021	18412	2.8	1.3	357341	0.8	0.5
3	adam-bomb-squad	0x7ab2...78c5	08/2021	25000	1.3	0.5	206478	0.5	0.2
4	autoglyphs	0xd4e4...7782	04/2019	512	2	0.7	206032	0.5	0.4
5	axie	0xf5b0...cb8d	04/2018	243000	1.6	0.3	285055	1.4	0.1
6	bored-ape-kennel-club	0xba30...5623	06/2021	10000	3.2	1.3	1423011	1.3	0.9
7	boredapeyachtclub	0xbc4c...f13d	04/2021	10000	4	1.4	11308109	1.4	1.6
8	collectvox	0xad9f...d34c	08/2021	8888	1.7	0.6	138121	0.4	0.5
9	cool-cats	0x1a92...050c	06/2021	9933	4	1.3	2299930	1.2	1.4
10	creature-world-collection	0xc92c...aafc	08/2021	10000	2.6	1	862273	0.9	0.9
11	cryptoadz-by-gremplin	0x1cb1...49c6	09/2021	7025	5.4	2.1	2964440	1.7	2.2
12	cryptokitties	0x0601...266d	01/2018	2009725	0.5	2.3	421288	1.6	0
13	cryptopunks	0xb47e...3bbb	01/2018	10000	10.2	4.3	53892061	2.4	3.7
14	cryptovoxels	0x7998...cf0c	06/2018	6210	1.5	0.4	54925	0.2	0.4
15	cyberkongz	0x57a2...4f37	04/2021	4147	3.1	1.6	824246	0.7	1
16	cyberkongz-vx	0x7ea3...7c8b	08/2021	14334	2.1	0.7	326717	0.5	0.4
17	deadfellaz	0x2aca...a17b	08/2021	10000	1.4	0.6	175900	0.6	0.5
18	decentraland	0xf87e...5d4d	04/2018	92598	9.2	9.7	3312118	7.5	0
19	doodles	0x8a90...992e	10/2021	10000	2.4	1.1	757788	0.8	0.7
20	fluf-world	0xccc4...a68d	08/2021	10000	2.8	1	434361	0.8	0.7
21	foundation	0x3b3e...5405	01/2021	103251	9.6	7.8	383853	11.2	0
22	galacticapes	0x12d2...4d14	09/2021	10000	3	1	540459	1	0.8
23	galaxyeggs	0xa081...7c48	09/2021	9999	2.8	1.1	577130	1.3	0.9
24	hashmasks	0xc2c7...6928	01/2021	16384	4.6	2.3	1493358	1.8	1.5
25	jungle-freaks-by-trosley	0x7e6b...4de0	10/2021	10000	3	1.4	887701	1.3	1.2
26	koala-intelligence-agency	0x3f5f...6360	08/2021	10000	2.5	1	393869	1	0.8
27	lazy-lions	0x8943...37e0	08/2021	10080	1.4	0.5	358509	0.6	0.5
28	lootproject	0xff9c...13d7	08/2021	7779	6.6	2.5	9458899	3.6	1.6
29	lostpoets	0xa720...f466	09/2021	27515	0.4	0.2	37669	0.1	0
30	meebits	0x7bd2...6bc7	05/2021	20000	2.1	0.8	1322740	0.6	0.3
31	mekaverse	0x9a53...ca8f	10/2021	8888	2.5	1.5	2173946	1.4	0.9
32	mutant-ape-yacht-club	0x60e4...a7c6	08/2021	30003	3.4	2	7253896	1.7	0.6
33	mutantcats	0xaaad...e46a	10/2021	10000	1.3	0.4	256031	0.5	0.4
34	pudgypenguins	0xbd35...2cf8	07/2021	8888	3.2	1.1	1764016	1.3	1.5
35	punks-comic	0x5ab2...c948	05/2021	10000	67	56.4	13674620	38.7	19.1
36	rari721	0x60f8...5ee5	05/2020	155346	6.4	5.7	2905148	14	0.1
37	rumble-kong-league	0xef01...909a	07/2021	10000	1.1	0.4	151521	0.5	0.2
38	sadgirlsbar	0x335e...b2d8	08/2021	10000	1.3	0.6	17836	0.4	0.4
39	sandbox	0x50f5...6d4a	12/2019	166464	0.9	0.5	200218	0.2	0.1
40	sneaky-vampire-syndicate	0x219b...2539	09/2021	8888	3.8	1.8	928516	1.4	1.2
41	somnium-space	0x595f...a0fa	10/2019	5025	1.3	0.5	244005	1.8	0.4
42	sorare	0x629a...6205	07/2019	329383	13.4	2.3	489009	0.7	0.7
43	supducks	0x3fe1...cbc5	07/2021	10000	2.8	1.1	544447	1	0.9
44	superrare1	0x41a3...850d	04/2018	4436	3.9	1.2	92844	0.2	0.6
45	superrare2	0xb932...b9e0	09/2019	4436	3.2	0.9	419872	0.7	2
46	the-doge-pound	0xf4ee...d043	07/2021	10000	2	0.8	725137	0.9	0.7
47	the-sevens-official	0xf497...187a	09/2021	7000	3.4	1.9	878788	2.1	1.3
48	thehumanoids	0x3a50...0edd	09/2021	10000	1.7	0.7	261715	0.7	0.6
49	tom-sachs-rockets	0x1159...5d26	08/2021	2000	76.9	60.9	16907477	58.6	54.3
50	veefriends	0xa3ae...beeb	05/2021	10255	1.8	1.4	662073	0.7	0.3
51	world-of-women-nft	0xe785...5330	07/2021	10000	1.6	0.5	539001	0.7	0.5
52	wrapped-mooncats	0x7c40...3572	03/2021	8903	15	6.1	981392	6.5	3.8




Figure C.5: Trades and unique trade partners. Each dot represents an address trading on NFT markets, positioned by the number of trades and unique trade partners. The size of the dot depicts the number of empirically identified suspicious trades.

Paper References

- [1] Arash Aloosh and Jiasun Li. Direct evidence of bitcoin wash trading. *SSRN Electronic Journal*, 2019. URL <https://papers.ssrn.com/abstract=3362153>.
- [2] Andreas Antonopoulos and Gavin Wood. *Mastering Ethereum*. O’Reilly Media, 2018.
- [3] Bloomberg. Jim chanos says nft market is rife with nefarious activity, 2021. URL <https://www.bloomberg.com/news/articles/2021-09-30/jim-chanos-says-nft-market-is-rife-with-nefarious-activity>.
- [4] Yi Cao, Yuhua Li, Sonya Coleman, Ammar Belatreche, and T.M. Mcginity. Detecting wash trade in financial market using digraphs and dynamic programming. *IEEE Conference on Computational Intelligence for Financial Engineering and Economics*, 2014.
- [5] Yi Cao, Yuhua Li, Sonya Coleman, Ammar Belatreche, and T.M. Mcginity. Detecting wash trade in financial market using digraphs and dynamic programming. *IEEE Transactions on Neural Networks and Learning Systems*, 2016.
- [6] Weili Chen, Tuo Zhang, Zhiguang Chen, Zibin Zheng, and Yutong Lu. Traveling the token world: A graph analysis of ethereum erc20 token ecosystem. *Proceedings of the World Wide Web Conference, WWW 2020*, pages 1411–1421, 2020.
- [7] Coindesk. The fast-growing nft market is problematic yet promising, 2021. URL <https://www.coindesk.com/business/2020/09/21/the-fast-growing-nft-market-is-problematic-yet-promising/>.
- [8] Lin Cong, Xi Li, Ke Tang, and Yang Yang. Crypto wash trading. *SSRN Electronic Journal*, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3530220.
- [9] Dipanjan Das, Priyanka Bose, Nicola Ruaro, Christopher Kruegel, and Giovanni Vigna. Understanding security issues in the nft ecosystem. *arXiv*, 2021. URL <http://arxiv.org/abs/2111.08893>.

- [10] Michael Dowling. Is non-fungible token pricing driven by cryptocurrencies? *SSRN Electronic Journal*, 2021. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3815093.
- [11] William Entriken, Dieter Shirley, Jacob Evans, and Nastassia Sachs. Erc-721 non-fungible token standard, 2018. URL <https://eips.ethereum.org/EIPS/eip-721>.
- [12] Mark Grinblatt and Matti Keloharju. Tax-loss trading and wash sales. *Journal of Financial Economics*, 71(1):51–76, 2004.
- [13] Martin Harrigan and Christoph Fretter. The unreasonable effectiveness of address clustering. *2016 International IEEE Conferences*, 2016.
- [14] Serkan Imisiker and Bedri Kamil Onur Tas. Wash trades as a stock market manipulation tool. *Journal of Behavioral and Experimental Finance*, 20:92–98, 2018.
- [15] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. An introduction to decentralized finance. *Complex Systems Informatics and Modeling Quarterly*, 2021.
- [16] Donald B. Johnson. Finding all the elementary circuits of a directed graph. *SIAM Journal on Computing*, 4(1):77–84, 1975.
- [17] Sashikanta Khuntia and Jamini Pattanayak. Adaptive market hypothesis and evolving predictability of bitcoin. *Economics Letters*, 2018.
- [18] Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. Mapping the nft revolution: Market trends, trade networks and visual features. *Scientific Reports*, 11(1), 2021.
- [19] Guéno le Le Penne c, Ingo Fiedler, and Lennart Ante. Wash trading at cryptocurrency exchanges. *Finance Research Letters*, 2021.
- [20] Ferdinand Regner, Andre Schweizer, and Nils Urbach. Nfts in practice - non-fungible tokens as core component of a blockchain-based event ticketing application. *40th International Conference on Information Systems, ICIS 2019*, 2019.
- [21] Robert Endre Tarajan. Depth first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1971.
- [22] Uniswap. Mint a new position uniswap v3, 2021. URL <https://docs.uniswap.org//protocol/guides/providing-liquidity/mint-a-position>.

- [23] Friedhelm Victor. Address clustering heuristics for ethereum. In *International Conference on Financial Cryptography and Data Security (FC 2020)*, pages 617–633, 2020.
- [24] Friedhelm Victor and Andrea Marie Weintraud. Detecting and quantifying wash trading on decentralized cryptocurrency exchanges. *Proceedings of the World Wide Web Conference, WWW 2021*, 2:23–32, 2021.
- [25] Victor von Wachter, Johannes Rude Jensen, and Omri Ross. Measuring asset composability as a proxy for ecosystem integration. *International Conference on Financial Cryptography and Data Security (FC 2021)*, 2021.
- [26] Qin Wang, Rujia Li, Qin Wang, and Shiping Chen. Non-fungible token (nft): Overview, evaluation, opportunities and challenges. *arXiv*, 2021. URL <http://arxiv.org/abs/2105.07447>.
- [27] Mark Weber, Giacomo Domeniconi, Jie Chen, Daniel Karl I. Weidele, Claudio Bellei, Tom Robinson, and Charles E. Leiserson. Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *CoRR*, 2019.

Blockchain-based Financial Infrastructure for Emerging Economies

Johannes Rude Jensen, University of Copenhagen, eToroX Labs

Victor von Wachter, University of Copenhagen

Omri Ross, University of Copenhagen, eToroX Labs

This paper has been published in the Proceedings of the European Conference on Information Systems (ECIS 2022).

Abstract The transformative capacity of blockchain technology is a frequently debated topic in the information systems (IS) and practitioner literature. Yet, rigorous and design-driven research remains relatively uncommon. We document an ongoing design process towards a blockchain-based IT-artifact comprising financial infrastructure for stakeholders in emerging economies. Working with a young NGO, we utilize the design-science research methodology (DSR) in the design, implementation, and evaluation of the IT-artifact. The artifact enables stakeholders to conduct basic financial services by computing transactions, maintaining a savings account, and receiving targeted stimulus payments. Following six months of iterative design and development, we released a global pilot version of the artifact. Over the first nine months, the pilot generated a dataset of 6.6 million transactions amongst 189,379 verified users. By conducting design-driven research, we contribute novel practical insights to the IS discourse on the transformative capacity of blockchain technology and information communication technologies (ICT) in emerging economies.

Keywords Blockchain, Design Science Research, Financial Infrastructure in Emerging Economies

Introduction

It is estimated that 1.7 billion people globally are unable to access the most basic financial services, a determining factor in preventing individuals in emerging economies from making the first leap out of poverty [4, 22]. The dominance of cash-based settlement procedures in emerging economies has been shown to impose ‘hidden tariffs’ on stakeholders throughout the value chain. Be it via storage and withdrawal costs, predatory middlemen, corruption, or other forms of financial exploitation, the hidden ‘cost of cash’ often levies significant financial penalties on unbanked individuals [3]. In this paper, we document interim results from an ongoing research project conducted in association with a young, non-governmental organization (NGO). We explore the feasibility of designing basic financial infrastructure for stakeholders in emerging economies, using blockchain technology. Because blockchain technology is an inherently transparent type of database architecture [7], we conjecture that this group of technologies may introduce a new level of transparency and accessibility to consumer-oriented financial services in emerging economies. We are utilizing the design-science research methodology (DSR) in the design, development, and evaluation of a blockchain-based IT-artifact. The artifact enables users to send financial transactions to each other and to receive stimulus payments from a shared pool of assets, directly to their digital wallet, which acts as a savings account. The research design is motivated by the research question: “How can blockchain technology support basic financial infrastructure for use in emerging economies?”. In this paper, we document the ongoing design process leading to the current iteration of the artifact consisting of six months of iterative design, development, and evaluation, followed by a nine-month pilot testing phase, in which the artifact was opened for use by a global group of stakeholders. The pilot version of the IT-artifact presented in this research-in-progress paper comprises a system of smart contracts deployed on Ethereum, a public, permissionless blockchain, and Fuse, a ‘side-chain’ enabling low-cost transaction processing and scaling. By exploring the capabilities and limitations of novel information communication technologies (ICT) through empirical and design-driven research, we seek to contribute to the growing body of IS literature on the transformational capacity of ICT.

Blockchain Technology in Emerging Economies

A blockchain is a replicated database, maintaining a shared state amongst a global network of nodes. Values, such as the balance of coins or tokens, are assigned to addresses and public keys, denoting the value in possession by the owner of the associated private key [7]. Nodes in the network compete to append the distributed database through a decentralized consensus mechanism. To elect the node in the network, tasked with propagating the

next block, either computationally hard problems or randomized selection is used [16]. The latest generations of blockchain technology have introduced the ability to deploy and execute basic scripting, known as ‘smart contracts’ in the shared database, paving the way for multiple interesting new forms of applications. To date, the most used applications emerged around financial services, providing innovative ways to conduct assets swaps or borrow money in decentralized money markets [13, 37].

In recent years, the practical and socioeconomic implications of blockchain technology have become a frequently discussed topic in the IS literature [18]. Scholars have proposed blockchain-based IT-artifacts for a wide range of problems in the financial industries, from shipping [23] to the settlement and clearing of financial transactions [29] and derivatives [5, 12] to issues concerning the management of sensitive data [6, 25] or ticketing [28]. Yet, the lion’s share of contributions to the IS blockchain literature approaches the technology from a theoretical angle [20, 30] examining how the unique properties of blockchain technology can create value for organizations [24], innovate business model designs [33] or operate in combination with other technologies [15].

The transformative capacity of ICT in emerging economies is broadly recognized within the IS discipline [21, 31]. Scholars have examined commercially driven financial infrastructure such as M-PESA in Kenya [1], PayTM in India [14], and KOMIDA in Indonesia [39]. Recent studies in the information systems genre and beyond portray an increasingly detailed picture of the challenges faced by individuals in emerging economies, identifying the high costs of banking and financial illiteracy as key drivers of financial exclusion [32]. To remedy these issues, it is argued, ICT artifacts must be mobile and scalable [27] while facilitating targeted payments to exposed minority groups [2]. To date, the majority of work on ICT in emerging economies explores the use of smartphones [35, 38] with little work done towards expanding our understanding of the transformational capacity of blockchain technology in emerging economies [19].

Methodology

The design process for the artifact is conducted in accordance with the DSR methodology [8]. As authors, we work alongside the NGO team members as active participants in the design and development process, while documenting the process throughout. At the time of the development of the present iteration of the artifact, the foundation team comprised an emerging markets economist, a project manager, and two software developers. The author team primarily contributes to the software development process through an in-depth understanding of blockchain technology and financial architecture. Prior to the release of the present iteration of the artifact, the authors worked alongside

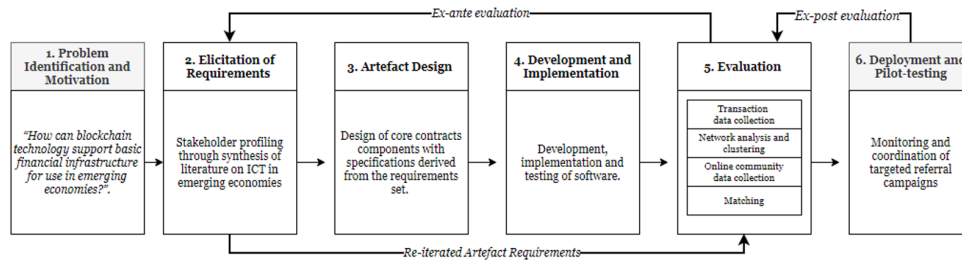


Figure D.1: The cyclical DSR workflow, exemplified by the activities in the latest cycle.

the foundation team members for a duration of six months. The day-to-day design and development workflow was conducted in smaller sprints leading into bi-weekly meetings in which progress within the general workflow was discussed.

Drawing on the rich literature on artifact evaluation practices, the project workflow was constructed with an emphasis on the iterative integration of cyclical evaluation results [10]. Figure D.1 displays the project workflow. The artifact design process was conducted with cyclical ex-ante stakeholder evaluations of the artifact requirements, alongside an ex-post evaluation cycle following the conclusion of the present iteration [36]. The evaluation process delineates a data-driven evaluation format in which granular blockchain transaction data is prepared for network analysis and cursively matched with publically available online data through non-participatory observation of users posting about the artifact on the social media pages Twitter and Facebook. The collection of qualitative data was conducted in accordance with the principles of netnography [17]. We observed online communities emerging around the artifact and gathered publicly available information, establishing context for the quantitative analysis.

Artifact Requirements for the Pilot Project The elicitation of effective artifact requirements [11] was subject to multiple cycles of ex-ante evaluation cycles, resulting in the version of the requirements presented in Table D.1 [36]. A target user profile was drawn from a synthesis of the extant literature on ICT and financial inclusion in emerging economies and injected into the requirements elicitation process [34]. The artifact requirements were initially targeted at a small user population, believed to be able to persuasively enroll users with similar needs [9].

Table D.1 The artifact requirements for the present iteration

	Requirement	Description
A	Low-cost transaction processing and stimulus distribution	(I) The artifact pilot must compute peer-to-peer transactions at a low-cost ratio to facilitate small transactions. (II) The artifact must facilitate the issuance of stimulus payments directly to active stakeholders' wallets.
B	Usability and accessibility	(I) The artifact must be accessible with low hardware requirements in order to embrace financial inclusion. (II) The artifact pilot must exhibit a capacity for persuasive enrollment and subsequent retention of stakeholders, through visible growth of the active stakeholder count, using only minor resources for dissemination and enrollment efforts.
C	Commercial viability	The artifact pilot should prove commercially viable through gradual integration into real (external) commercial activities.

Artifact Demonstration and the Global Pilot Design

The current iteration of the artifact comprises a set of six smart contracts, deployed on the Ethereum blockchain and the Fuse sidechain in parallel (Figure D.2). The Ethereum blockchain was chosen due to the high level of security and wealth of developer tooling available. The Fuse sidechain is a replicated version of the Ethereum blockchain, utilizing a ‘delegated proof of stake mechanism’. The Fuse sidechain requires fewer servers on the network, which facilitates fast low-cost transaction processing¹. In comparison, the virtual machines of both blockchains are fairly similar, whereas the Fuse sidechain sacrifices decentralization in order to gain a scalability advantage over Ethereum. We balance the lower level of decentralization on the Fuse blockchain by deploying the contracts with the largest security requirements on the Ethereum blockchain, whilst deploying the contracts with a lot of transactions on the Fuse blockchain. Because the Fuse blockchain is a replicated low-cost environment to Ethereum, the smart contracts can communicate through a ‘bridge-contract’. A simple browser-based user interface (UI) was designed for browser and mobile accessibility. Through the UI the user can compute (I) transactions (send and receive) of tokens, and (II) ‘claims’ of basic income stipends in tokens issued in cyclical intervals. The contract

¹<https://docs.fuse.io/the-fuse-chain/overview>, accessed Jul 2021

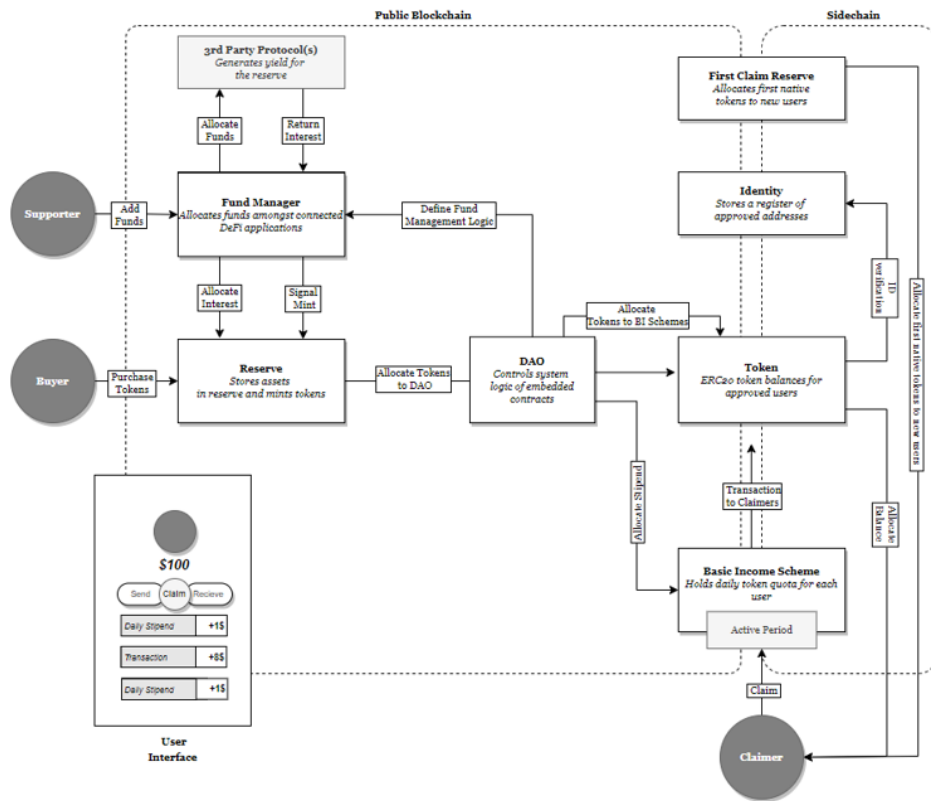


Figure D.2: Artifact overview and critical flows between smart contracts

system taxonomy is constructed as follows: The token is the primary unit of exchange for the artifact. It maintains the balance of all user’s addresses and acts as a claim on the yield generated by the crypto assets in the Reserve contract. Computing transactions through the UI communicates directly with the Token contract interface. Prior to computing any transaction, the Token checks the recipient’s address with the Identity contract, which stores a list of approved addresses. To become approved, users must execute a signup verification flow utilizing their email address, through which a single token is issued to the address through the First Claim Reserve contract.

Supporters can contribute crypto assets to the Fund Manager contract which allocates funds to third-party applications, generating interest yield through lending the donated crypto assets on a money market². The yield generated by the Fund Manager is sent to the Reserve contract, where tokens

²For the current iteration of the artifact, the Fund Manager allocates funds to ‘Compound’ a smart contract-based money market where borrowers of crypto assets pay a variable interest rate to lenders. The interests generated by the Fund Manager are submitted to the Reserve. We invite the reader to view ‘Compound’ on compound.finance.

are minted at a free-floating ratio to the USD value of the assets in the reserve, approximating an exchange rate of \$0.0001 per token. Buyers, who wish to purchase Tokens, can do so directly from the Reserve contract by sending a transaction in any approved crypto asset to the Reserve. The Reserve contract then allocates Tokens to the DAO contract, which submits a stipulated amount to the Basic Income contract. The Basic Income contract stipulates an ‘active period’ in which a given amount of tokens is allocated for claiming by verified users. Users call the ‘claim’ function in the Basic Income contract through the UI, triggering a transaction of the daily stipend to their accounts. For the nine-month pilot, stakeholders were able to ‘claim’ tokens from the Basic Income contract at a daily cadence. Additional allocations of the tokens were made to referral contracts through which stakeholders could refer friends for an additional sign-up bonus.

Artifact Evaluation

We follow a data-driven evaluation format in which blockchain transaction data is prepared for network analysis and checked against qualitative user data obtained through netnographic observational studies of user behavior in online communities [36]. In Table D.2 we summarize the results of the evaluation process by mapping the artifact requirements against the observed behavior.

The current iteration of the artifact was deployed on the Ethereum and Fuse networks and opened to use on 01.08.2020. Over the course of nine months, the artifact processed 6.6 million transactions amongst 189,379 verified individuals. The mean transaction fee stood at \$0.00001, or 1/1000 of a cent, per transaction with an average processing time of 2.5 seconds. The average transaction value is \$0.045, with a total of 151,858 participants computing three or more transactions. Through the course of the pilot, 44 donations were made to the Fund Manager. The contract currently holds crypto assets valued at about \$120,000 which yielded an average return of 5.56% p.a. for the duration of the pilot. We explored the dataset utilizing a graph-based network analysis, in the open-source network analysis tool Gephi, creating a network graph with a randomized sample ratio of 1:54 of the total data set. Nodes represent individual addresses, colored by their number of connected edges. Edges are weighted and represent transaction volumes (Figure D.2).

As evident, most transactions on the network are computed by claimers (6.3 million), withdrawing a daily stipend from the Basic Income contract, indicating a dominant ‘claim and hold’ pattern where stakeholders sign up and use the artifact to claim the tokens successively (with some stakeholders having claimed up to 250 times). This pattern is partially confirmed by the large number of stakeholders claiming from several generations of the Basic

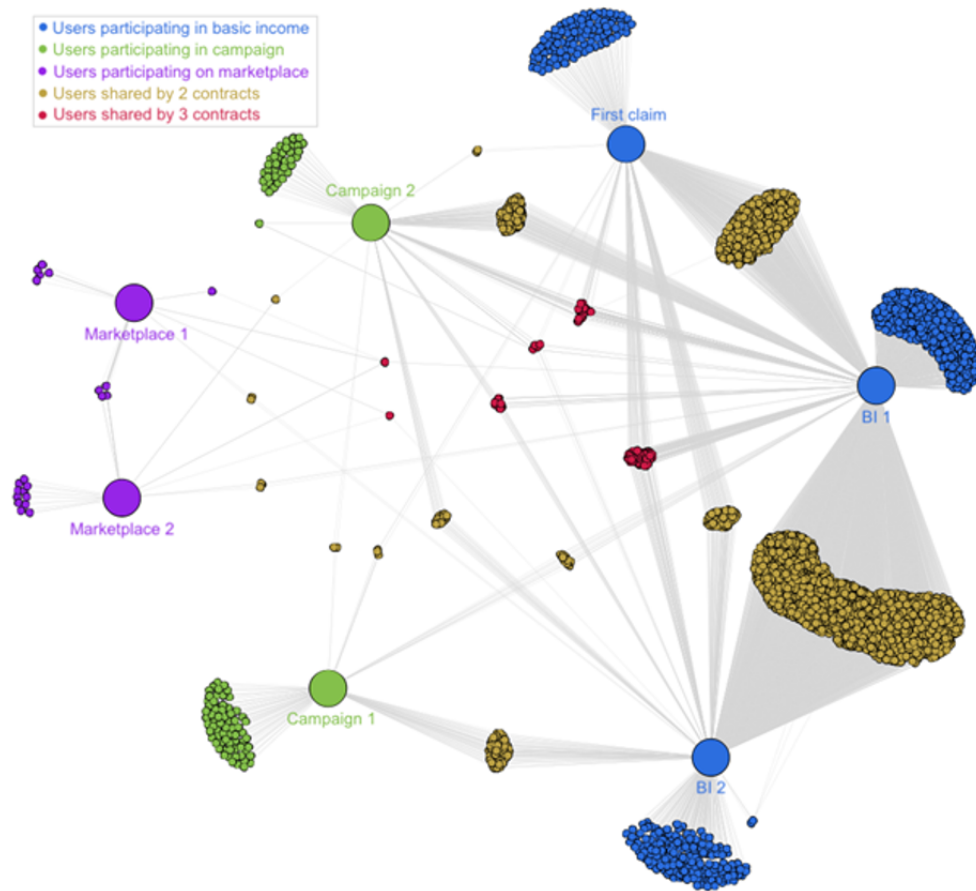


Figure D.3: Network analysis of transactions

Income contracts (BI1-BI2)³

Throughout the nine-month pilot phase, we noted the appearance of several interesting transaction clusters. The qualitative data collected through the netnographic method linked these clusters to organically emerging marketplaces, in which stakeholders traded items, services, or other crypto assets against the token on decentralized exchanges (purple nodes).

This revealed several fascinating stories warranting deeper empirical study of user behavior, including (I) a group of stakeholders organizing through Facebook to build an online portal for the sale and listing of used items traded with the artifact Token (Marketplace 1), and (II) a liquid market for the token on a pre-existing decentralized exchange where stakeholders appear to be trading the token against other crypto assets, outside of the minting price range

³Please Note: Due to a technical issue, the Basic Income contract was redeployed midway through the pilot (BI1-BI2).

(Marketplace 2). The most recent evaluation cycle was completed following the recent conclusion of the pilot phase. Table D.2 summarizes these results.

Discussion and Interim Conclusions

While it is not yet clear whether the group of technologies associated with the term ‘blockchain technology’ in the IS literature, is the appropriate design choice for financial infrastructure in emerging economies, the ongoing work presented here provides guidance on the feasibility of implementing a basic artifact for the emerging markets context. The ongoing work presented here is guided by the broadly posed research question: “How can blockchain technology support basic financial infrastructure for use in emerging economies?”. Through the design, development, and evaluation of a blockchain-based IT-artifact, we demonstrate the feasibility of implementing a modest digital wallet with the ability to process transactions and distribute basic-income stimulus payments to stakeholders.

In the ongoing pursuit of the research question, we define three key product requirements in collaboration with a small team of stakeholders at an NGO, drawing on the DSR methodology. The product requirements were designed in collaboration with the NGO design partners and delineated the ideal properties for future iterations of this artifact. We find that an ideal iteration of the artifact must (I) process transactions with a low-cost base to facilitate direct stimulus or universal basic income payments to wallets (II) be accessible for use on low-cost smartphones with an internet connection and (III) facilitate commercial transactions between a large set of users, facilitating real economic value creation. We document the current progress on achieving these requirements through the deployment of a nine-month prototype. The prototype utilizes blockchain side-chain architecture as a scaling mechanism, conducting over 6.6 million transactions between 189,379 verified individuals across multiple emerging economies over the course of nine months. The present iteration of the artifact utilizes a test deployment of a UBI scheme by which stakeholders can withdraw tokens minted against a claim in a reserve. This appears to have introduced strong growth incentives resulting in relatively fast adoption of the prototype amongst a wide userbase, however, it does not appear to have generated a large amount of authentic economic transactions between the stakeholders at this point.

Given the ongoing nature of the design process for the iteration of the artifact presented in this preliminary paper, there are multiple limitations to the present study. Primarily, the design risk for the viability of the present iteration of the artifact is social and user-oriented [36]. While the artifact is designed for organic growth through commercial adoption in emerging economies, the present study does not address or document the user experience of stakeholders in emerging economies. As a consequence, the interpre-

Table D.2 The artifact requirements for the present iteration

	Requirement	Description
A	Low-cost transaction processing and stimulus distribution	The requirement was addressed: The side-chain architecture has proven efficient in computing transactions at a cost ratio far below the costs on the Ethereum blockchain. The average transaction cost does not appear to have restrained the distribution of basic income payments to users' digital wallets.
B	Usability and accessibility	The requirement was addressed: The artifact is accessible with an internet connection via regular browsers and supports lower-end mobiles such as featurephones. Furthermore, the surprisingly rapid growth of stakeholders using the pilot version of the artifact to claim a daily stipend and the popularity of the referral campaigns is indicative of the capacity of organic growth. We interpret the prevalence of users claiming regularly as indicative that users are reasonably comfortable using the artifact.
C	Commercial viability	The requirement was insufficiently addressed: The data set revealed interesting activities conducive to minor commercial activities settled in the artifact. Yet, the prevalence of the passive 'claim-and-hold' pattern, is indicative that a large number of passive stakeholders appear to be claiming the token with the hope that it will appreciate in value over time. This pattern does not make any clear indication of the feasibility that the artifact can support commercial processes in an emerging economy. Evaluating the commercial viability of the artifact is likely to require on-the-ground empirical studies of stakeholder engagement with the artifact.

tation of the provisional findings presented in this paper is liable to Type I errors, known as false positives [26].

While the artifact provides an indication of the feasibility of implementing financial infrastructure with blockchain technology, extensive empirical trials in emerging economies are a requisite, if a contribution to the IS literature on financial inclusion is to be made. Empirical trials must emphasize the path to adoption by linking transactional data to behavioral and observational data, contributing a new level of granularity in the collection of empirical data on emerging economies [21]. Future work on this artifact will include an expansion of the research design with a comparable study in which the artifact is benchmarked with existing financial services in emerging economies. Additionally, we intend to conduct a variety of improvements and enhancements whilst designing additional empirical trials to better gauge user behavior. For the prototype presented here, we omit discussion on certain crucial qualifiers for financial infrastructure. Given the pseudonymous nature of accounts on Ethereum, privacy features are a necessary requirement for real-world success.

Paper References

- [1] Stephen Agyepong and Hossana Twinomurinzi. Facilitating Financial Inclusion using ICT: Lessons from M-PESA and E-ZWICH. *European Conference on Information Systems (ECIS 2016)*, 2016.
- [2] Jenny Aker, Rachid Boumnijel, Amanda McClelland, and Niall Tierney. Payment Mechanisms and Antipoverty Programs: Evidence from a Mobile Money Cash Transfer Experiment in Niger. *Economic Development and Cultural Change*, 65(1), 2016.
- [3] Bhaskar Chakravorti. The hidden Costs of Cash. *Harvard Business Review*, 18, 2015.
- [4] Asli Demirguc-Kunt, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess. *The Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution*. World Bank Reports, 2018.
- [5] Benjamin Egelund-Müller, Martin Elsman, Fritz Henglein, and Omri Ross. Automated Execution of Financial Contracts on Blockchains. *Business and Information Systems Engineering*, 59(6):457–467, dec 2017.
- [6] Benedict Faber, Georg Cappelen Michelet, Niklas Weidmann, Raghava Rao Mukkamala, and Ravi Vatrapu. BPDIMS: A Blockchain-based Personal Data and Identity Management System. *Hawaii International Conference on System Sciences (HICSS 2019)*, 6:6855–6864, 2019.
- [7] Florian Glaser. Pervasive Decentralisation of Digital Infrastructures: A Framework for Blockchain enabled System and Use Case Analysis. In *Hawaii International Conference on System Sciences (HICSS 2017)*, 2017.
- [8] Shirley Gregor and Alan R Hevner. Positioning and Presenting Design Science Research with Maximum Impact. *MIS Quarterly*, 37(2):337–355, 2013.
- [9] Ole Hanseth and Kalle Lyytinen. Design Theory for Dynamic Complexity in Information Infrastructures: The Case of Building Internet. *Journal of Information Technology*, 25(1):1–19, 2010.

- [10] Marlien Herselman and Adele Botha. Evaluating an Artifact in Design Science Research. *ACM International Conference Proceeding Series*, 2015.
- [11] Ann M Hickey, Alan M Davis, and Tony O’Hagan. A Unified Model of Requirements Elicitation. *Journal of Management Information Systems*, 20(4):65–84, 2004.
- [12] Johannes Rude Jensen and Omri Ross. Settlement with Distributed Ledger Technology. In *International Conference on Information Systems (ICIS 2020)*, 2020.
- [13] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. An Introduction to Decentralized Finance (DeFi). *Complex Systems Informatics and Modeling Quarterly*, 26(3):46–54, 2021.
- [14] Tanmay Joshi, Sharmistha Swasti Gupta, and Nimmi Rangaswamy. Digital Wallets ’Turning a Corner’ for Financial Inclusion: A Study of Everyday PayTM Practices in India. *IFIP Advances in Information and Communication Technology*, 552:280–293, 2019.
- [15] Erik Karger. Combining Blockchain and Artificial Intelligence - Literature Review and State of the Art. *International Conference on Information Systems (ICIS 2020)*, pages 1–17, 2020.
- [16] John Kolb, Moustafa AbdelBaky, Randy H Katz, and David E Culler. Core Concepts, Challenges, and Future Directions in Blockchain. *ACM Computing Surveys*, 53(1):1–39, 2020.
- [17] Robert V. Kozinets, Pierre-Yann Dolbec, and Amanda Earley. Netnographic Analysis: Understanding Culture through Social Media Data. *The SAGE Handbook of Qualitative Data Analysis*, pages 262–276, 2014.
- [18] Olga Labazova. Towards a Framework for Evaluation of Blockchain Implementations. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [19] Guillermo Jesús Larios-Hernández. Blockchain Entrepreneurship Opportunity in the Practices of the Unbanked. *Business Horizons*, 60(6):865–874, 2017.
- [20] Juho Lindman, Virpi Kristiina Tuunainen, and Matti Rossi. Opportunities and Risks of Blockchain Technologies: A Research Agenda. In *Hawaii International Conference on System Sciences (HICSS 2017)*. Hawaii International Conference on System Sciences, 2017.

- [21] Edda Tandi Lwoga and Raphael Zozimus Sangeda. ICTs and Development in developing Countries: A Systematic Review of Reviews. *The Electronic Journal of Information Systems in Developing Countries*, 85 (1), 2019.
- [22] Lakshmi Mohan and Devendra Potnis. Real-time decision-making to serve the Unbanked Poor in the Developing World. *SIGMIS-CPR 2017 - Proceedings of the 2017 ACM SIGMIS Conference on Computers and People Research*, pages 183–184, 2017.
- [23] Kristoffer Nærland, Christoph Müller-Bloch, Roman Beck, and Søren Palmund. Blockchain to rule the Waves - Nascent Design Principles for reducing Risk and uncertainty in decentralized Environments. In *International Conference on Information Systems (ICIS 2017)*, 2017.
- [24] Nadine Kathrin Ostern, Michael Rosemann, and Jürgen Moormann. Determining the idiosyncrasy of Blockchain: An Affordances Perspective. *International Conference on Information Systems (ICIS 2020)*, 2020.
- [25] José Parra Moyano and Omri Ross. KYC Optimization Using Distributed Ledger Technology. *Business and Information Systems Engineering*, 59 (6):411–423, dec 2017.
- [26] Jan Pries-Heje, Richard Baskerville, and John Venable. Soft Design Science Research: Extending the Boundaries of Evaluation in Design Science Research. In *Proceedings from the 2nd International Conference on Design Science Research in IT (DESRIST)*, pages 18–38, 2007.
- [27] Carol Realini and Karl Mehta. *Financial Inclusion at the Bottom of the Pyramid*. FriesenPress, San Francisco, US, 2015.
- [28] Ferdinand Regner, André Schweizer, and Nils Urbach. NFTs in Practice - Non-fungible Tokens as core component of a Blockchain-based Event Ticketing Application. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [29] Omri Ross, Johannes Rude Jensen, and Truls Asheim. Assets under Tokenization: Can Blockchain Technology Improve Post-Trade Processing? In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [30] Matti Rossi, Christoph Mueller-Bloch, Jason Bennett Thatcher, and Roman Beck. Blockchain Research in Information Systems: Current Trends and an Inclusive Future Research Agenda. *Journal of the Association for Information Systems*, 20(9):1388–1403, 2019.
- [31] Narcyz Roztocky and H Roland Weistroffer. IS/IT in Developing and Emerging Economies. In *AMCIS 2007*, 2007.

- [32] Sebastian Schuetz and Viswanath Venkatesh. Blockchain, Adoption, and Financial Inclusion in India: Research Opportunities. *International Journal of Information Management*, 52, 2020.
- [33] Stefan Seebacher and Maria Maleshkova. A Model-driven Approach for the Description of Blockchain Business Networks. *Hawaii International Conference on System Sciences (HICSS 2018)*, 9, 2018.
- [34] Maung K Sein, Ola Henfridsson, Sandeep Purao, Matti Rossi, and Rikard Lindgren. Action Design Research. *MIS Quarterly*, 9(2):37–56, 2011.
- [35] P. K. Senyo and Ellis L.C. Osabutey. Unearthing Antecedents to Financial Inclusion through FinTech Innovations. *Technovation*, 98, 2020.
- [36] John Venable, Jan Pries-Heje, and Richard Baskerville. FEDS: a Framework for Evaluation in Design Science Research. *European Journal of Information Systems*, 25(1):77–89, 2016.
- [37] Victor Von Wachter, Johannes Jensen, and Omri Ross. Measuring Asset Composability as a Proxy for Ecosystem Integration. In *International Conference on Financial Cryptography and Data Security (FC 2021). International Workshops.*, 2021.
- [38] Bruno L. Yawe and Jaideep Prabhu. Innovation and Financial Inclusion: A Review of the Literature. *Journal of Payments Strategy and Systems*, 9(3):215–228, 2015.
- [39] Adrian Yeow and W.K. Lim. KOMIDA: Making Microfinance Digital in Indonesia. In *International Conference on Information Systems (ICIS 2018)*, 2018.

Kickstarting Blockchain: Designing Blockchain-based Tokens for Equity Crowdfunding

Tobias Guggenberger, University of Bayreuth
Benjamin Schellinger, University of Bayreuth
Victor von Wachter, University of Copenhagen
Nils Urbach, Frankfurt University of Applied Sciences

This paper has been published in *Electronic Commerce Research Journal* (2023).

Abstract Blockchain-based tokens seek to overcome the friction and opaqueness of the legacy financial infrastructure in the company funding process, particularly in the early-stage and equity crowdfunding domain. While Initial Coin Offerings and Security Token Offerings proposed a solution for crowdfunding, early-stage companies still face challenges in using blockchain as an alternative equity funding infrastructure. In this context, the idea of blockchain-based equity tokens remains hypothetical. In addition, the literature lacks a design theory for the development and implementation of blockchain-based equity tokens. This research bridges this gap by designing, developing, and evaluating an equity token prototype for crowdfunding, following the design science research approach. We propose a refined crowdfunding model and derive seven design principles that contribute to the design theory of equity tokens. The research results show that blockchain-based equity tokens improve efficiency, transparency, and interoperability while meeting regulatory requirements and facilitating secondary market trading.

Keywords Blockchain, Design science, Equity crowdfunding, Initial coin offering, Security token offering, Tokens

Introduction

Entrepreneurship is a desirable goal for economies to foster innovation, stimulate economic growth, and create employment [24, 11, 66]. During the early stages of entrepreneurship, funding is often indispensable to drive forward and implement an idea or a project. Therefore, funding as a method of raising capital outside of operating cash flow is of utmost importance to mitigate early-stage companies' operational risks and secure long-term growth. However, entrepreneurs still face various problems during and upon a traditional early-stage funding process, including geographical constraints, exclusive networks, and the involvement of multiple intermediaries [24, 20, 23]. In addition, it is slow and expensive owing to the plethora of intermediaries involved [20, 13, 31]. In an endeavor to improve early-stage funding, equity crowdfunding emerged as an alternative funding tool, reaching a total funding amount of over \$1.5bn globally in 2018 [10]. Equity crowdfunding is a crowd-based form of issuing company shares in exchange for capital via an Internet platform giving investors equity-like rights. These rights make equity crowdfunding more similar to the issuance of shares than they mimic the idea of donation- or reward-based crowdfunding [37, 51, 22]. Although equity crowdfunding optimizes prior forms of early-stage funding, it lacks broad liquidity, entails bureaucracy and high administrative costs while still relying on trusted intermediaries, such as centralized platform providers [81, 8, 63]. Initial Coin Offering (ICO) via blockchain technology proposed an alternative approach to traditional crowdfunding and enabled more efficient crowdfunding processes, thus, democratizing early-stage investments [37, 13]. In an ICO, investors generally trade in their cryptocurrency in exchange for a utility token, representing the right to use a particular offered service [12, 34]. Following substantial growth in 2017 (\$6.2 bn) and 2018 (\$7.8 bn), total funds raised through ICOs decreased to \$0.3 bn in 2019 [40]. Consequently, initial enthusiasm has turned into declining investment in ICOs, mainly because of unclear regulation, limited configurability, and insufficient investor protection [44]. The stagnant technological improvement of the traditional funding process and the lack of regulatory compliance of ICOs led to the latest development of Security Token Offering (STO). A security token is a digital representation of particular security issued and managed on a blockchain using smart contracts and computer code that executes arbitrary business logic [75, 34]. Unlike utility tokens, security tokens issued via STOs comply with regulatory requirements by default, grant the token holder an underlying value, and, eventually, present a more mature form of token sales [49, 50, 42]. As such, STOs can be seen as an alternative to equity crowdfunding platforms. Thus, we state that blockchain technology improves the efficiency, transparency, and interoperability of conventional equity crowdfunding. In addition, the configurability of smart contracts allows regulatory compliance and creates liquidity, facilitating trading in the secondary market. Even though researchers recog-

nize the value of blockchain for equity crowdfunding, theory in this area is limited [36]. In sum, existing research [37, 81, 1] focuses on the potential of blockchain for equity crowdfunding but lacks design knowledge in this context. However, design theory is a prerequisite to understanding how such systems should be implemented and effectively foster added value [62, 6]. To address this gap, we define the following research questions:

RQ: How can blockchain be incorporated as an alternative infrastructure for equity crowdfunding?

Our research objective is to bridge the identified gap in the IS literature and answer the question by designing, implementing, and evaluating a blockchain-based equity token prototype following the design science research (DSR) paradigm [32, 38, 48]. In doing so, we aim to respond to Treiblmaier et al. [76] call to design a security token and explore its potential to reduce information asymmetries, improve operations, and ultimately allocate capital more efficiently. In addition, we take up the research agenda by Kranz et al. [42] and the call of Perdana et al. [59] and focus on a particular security token, i.e., an equity token. This paper is the first to design a blockchain-based equity token for crowdfunding to the best of our knowledge. Overall, we seek to make the following primary contributions. First, developing a blockchain prototype will allow us to gain practical insights into the opportunities and challenges of implementing complex blockchain-based solutions, expanding the blockchain-based equity token research, and the early-stage funding fields. Second, we seek to deepen the understanding of mandatory requirements and the infinite design space of blockchain-based equity tokens, contributing to design theory in this field by developing and evaluating an instantiation of a blockchain-based equity token for crowdfunding. Third, we extend the crowdfunding model developed by Haas et al. [33] by outsourcing traditional financial and operational services to smart contracts and adding new stakeholders. Fourth, we seek to derive seven generalized design principles (DP) to guide the design and development of blockchain-based equity tokens. The remainder of this paper is structured as follows: In Section 2, we present the principles of traditional early-stage funding and equity crowdfunding, followed by blockchain-based crowdfunding. Next, in Section 3, we present our DSR approach, while in Section 4, we elaborate on the instance problem, i.e., equity crowdfunding. Section 5 shows the derived software requirements and provides a detailed account of the software prototype development. In Section 6, we evaluate the prototype and the research approach. Section 7 generalizes and discusses the results based on both the literature and semi-structured interviews and derives design principles. We conclude with a summary, highlighting limitations and outlining future research directions in Section 8.

Background

Early-stage Funding and Equity Crowdfunding

Early-stage funding Entrepreneurship is a pursuable goal in every economy as literature has long identified the role of entrepreneurship in enhancing innovation, economic growth, and job creation [24, 11, 66]. When looking to thrive on an idea or project, early-stage entrepreneurial funding is often inevitable. However, due to the short business history, funding instruments like loans or bonds provided by financial institutions or other market participants are not available [31, 77]. Thus, the financing of early-stage companies takes place in the private market through the issuance of large investment tickets, which excludes small investors from participating in these companies. Consequently, this led to the establishment of an inaccessible and concentrated market for early-stage funding with specialized participants [18]. In particular, specialized intermediaries, which are reputed to be experienced with high uncertainty and principal-agent problems in entrepreneurial financing, serve the market [31, 17]. In this context, the US-style venture capital process has been subject to criticism ever since and is regarded as one of the major constraints for the full exploitation of the economic potential of entrepreneurship [24, 20, 23]. The process of entrepreneurial funding takes a substantial amount of time, involves many different parties, leads to cumbersome bureaucracy regarding the preparation of contracts, and requires sound knowledge and a personal network in the industry. In addition, it is slow and expensive owing to the plethora of intermediaries involved [20, 13, 31, 70]. Consequently, this stagnant funding process led entrepreneurs to look for ways to improve the traditional venture capital funding system [7].

Equity crowdfunding Equity crowdfunding platforms are a promising improvement heavily discussed in the literature [51, 22]. Equity crowdfunding is a crowd-based form of issuing company shares in exchange for capital via an Internet platform [22]. Websites usually host these platforms, while web-based software often facilitates interaction between entrepreneurs and investors willing to fund their projects [77]. While in the traditional system, money is provided towards selected projects, crowdfunding can be accessed by a larger group that decides to invest a smaller contribution into a potentially successful company [70]. For example, EquityNet offers companies a platform to promote their venture, including business cases and financial figures. The investment in a company is a stark contrast to well-known fundraising platforms like Kickstarter and GoFundMe, which are raising money for a project without expectation of return (i.e., they are in contrast to donation-based or reward-based for non-monetary rewards) [37, 51]. Both conventional and equity crowdfunding share common characteristics: Early and global access via an Internet platform makes it possible to gather a contributing community

around the company from the very beginning. Therefore, these crowdfunding mechanisms facilitate the attraction of investors, create a brand, and increase media coverage [37, 77]. Yet, crowd interest is often more diverse and involves social intent [77], and crowdfunding investments are spread across a broader range of companies than traditional venture capital. But whereas Kickstarter has revolutionized the fundraising space for reward-based projects, the adoption of equity crowdfunding platforms is still limited [12]. In sum, our literature analysis reveals that there is no overall satisfying funding mechanism to answer the specific needs of early-stage companies in a fast, affordable, and equal manner. Thus, we explore a novel blockchain-based funding mechanism that tries to address the shortfalls to bring equal benefits to entrepreneurs and investors.

Blockchain-based Crowdfunding

Blockchain The interest of academia and practice in blockchain technology first arose after the Bitcoin white paper by Nakamoto [52], who proposed a peer-to-peer (P2P) digital currency. Many researchers and practitioners state that blockchain can radically change an extensive range of business processes [52, 55, 62]. Blockchain describes a distributed ledger that records and secures transactions in a decentralized network [62]. A trust-free consensus algorithm, run by the participating nodes, determines the order of all executed transactions and the currently valid blockchain state [30].¹ In addition, blockchain describes an algorithmic protocol with the potential for global disintermediation through the decentralization of transaction confirmation between participants who previously did not trust one another [72]. With its decentralized application platform, using a virtual machine (EVM) and a built-in Turing-complete programming language, the Ethereum blockchain facilitates the use of smart contracts [9]. Smart contracts describe an algorithmic transaction protocol that automatically executes the terms of a contract on a blockchain to achieve trust between two or more unacquainted participants [73]. The consensus protocol ensures the enforcement of these scripts and can reduce transaction costs and improve settlement speed [6, 30, 73].

Blockchain tokens and distribution A token is a series of characters that identifies a specific asset right or asset class [71]. Technically tokens can be used in several cases, e.g., in an internal unit of account, in the facilitation of transactions, or to grant token holders certain types of privileged access [9, 71]. While a native token is deeply implemented on the blockchain protocol (e.g., Bitcoin or Ether), tokens issued on top of the blockchain layer are usually managed by smart contracts [39, 34]. Since the Ethereum blockchain was the

¹A consensus algorithm is only purely considered trust-free if it does not rely on trusted validating nodes, e.g., in the context of a private blockchain.

Table E.1 ERC Token Standards on the Ethereum Blockchain.

Token type	Fungible	Non-fungible	Multiple	Security-token
Characteristics	Divisible	Unique	Divisible and unique	Regulatorily compliant
Use cases	Currencies, access or voting rights	Collectibles, tickets, digital artwork	Equity, real estate, game items	Financial securities

first to allow for implementing business logic using smart contracts, different standards of the token interface have emerged over the years to ensure interoperability on the platform. The Ethereum community, developer, and token holders can propose improvements (EIP, Ethereum Improvement Proposals) on smart contract functionalities, resulting in the relevant Ethereum Request for Comments (ERC), such as ERC20, ERC721, ERC1155, and EIP1400 (see Table E.1). Chiefly, tokens can be divided into utility and security tokens. Utility tokens are issued via Initial Coin Offerings (ICOs) and provide access or payment to digital services, granting the issuing company complete control over which rights and claims are connected to the token [12, 34]. The literature confirms the benefits of ICOs as a funding alternative over traditional crowdfunding methods [29, 3] and extensively analyzes its success factors [14, 27, 46, 57]. However, ICO tokens also have drawbacks that negatively affect the use of the platform. Although the flexibility can explain the previous dominance of utility tokens, the issuing company, regulatory loopholes, a broad investing community, and the efficiency of blockchain [15], the majority of ICOs may have been misguided or even fraudulent with no intention of fulfilling the project pipeline [43]. Concerns have been raised about the lack of regulatory compliance and basic investor protections, as ICO tokens are considered securities in disguise, owing to their reward-based character [44]. In addition, there is a lack of incorporating real-world security regulations on the blockchain and supervising mechanisms steering the company [79]. Consequently, ICO success is bound to the attractiveness of the underlying value, e.g., the company and the granted token rights. However, often the token issued does not inhibit rights and thus has no underlying value. Recently, the advancement of ICOs to security token offerings (STOs) holds new promises for token-based funding [50]. Unlike ICOs, STOs cater for the whole funding lifecycle, i.e., issuance, maintenance, dissolution, regular communication (e.g., quarterly reporting), voting rights, and equity-specific transactions (e.g., dividends). In addition, STOs apply to cross-border regulation with on-chain and off-chain interactions by design using programmable smart contracts and hence present a more mature form of token sales [49, 42]. Security tokens represent tokenized ownership, i.e., a digital representation thereof, and are

subject to security regulation [75, 34]. Equity tokens are a subclass of security tokens and represent ownership of equity that entails rights and obligations under equity legislation, e.g., the right to dividends or voting rights. Thus, equity tokens are digital representations of shares on a blockchain [34]. On the other hand, a vast number of decentralized finance (DeFi) projects, such as Uniswap, Aave, or Curve, primarily emulate ownership by issuing governance tokens. However, these governance tokens only grant utility token-like rights to these DeFi protocols, i.e., voting rights in project development, and thus do not represent a regulated form of a security, or more specifically, an equity token for these projects [65]. Both utility and security tokens are fungible and tradable, but their value is derived differently from the underlying asset or service they represent [75]. Due to the infinite design options and legal complexity, it is not easy to classify tokens, and in fact, many tokens are between the categories of utility and security. If a token is either a utility or security is commonly tested by a legal precedent determining security status. In this context, the Security Exchange Commission (SEC) in the U.S. has developed the Howey test to assess whether a token can be classified as a security and thus needs to be regulated. The SEC Howey Test has evolved as a de facto simplifying standard within the blockchain community once a token is considered a security. According to the test, a token will be classified as security if all four of the following requirements are fulfilled: (i) investment of money, (ii) common enterprise, (iii) profit expectation, and (iv) solely on the effort of others. The legal status of utility tokens is surrounded by controversy due to the grey area of their true economic value. Accordingly, regulation across the globe has been different, ranging from pending regulation to promotion on a case-by-case evaluation to outright ban [43]. Security tokens go along with a more expensive initial registration, more obligations to investors during the lifecycle of the security, and potential fines if investor rights are not met [15]. In what follows, we take the U.S. law as our starting point and therefore cannot ensure that it applies to early-stage companies in other jurisdictions. An early-stage company could circumvent traditional equity funding vehicles like venture capital or private equity by issuing equity tokens through blockchain. The token issuance process purely relies on P2P mechanisms instead of the matchmaking process by crowdfunding platforms and banks between campaign creators and potential investors [33, 67]. Unlike conventional crowdfunding, token sales offer common advantages that make it more attractive to global investors. There is a deeper pool of liquidity, and ownership becomes divisible and thus tradable [2]. Companies can develop their proprietary blockchain protocol to issue and sell native tokens [29] or use existing infrastructure, e.g., the Ethereum blockchain, and sell on-chain utility tokens [2, 14]. Even though researchers recognize the merits of token sales, the literature on blockchain-based crowdfunding is limited. Arifin et al. [1] propose that blockchain-based crowdfunding can leverage financial inclusion and reduce challenges associated with platform operators. Zhu and

Zhou [81] analyze blockchain-based equity crowdfunding in China and find that blockchain can reduce friction, thus facilitating the circulation of equity shares. In addition, blockchain enables P2P transactions, improves governance, and provides regulators with necessary market information [81]. In a Delphi Study, Heieck [37] confirmed driving the merits of blockchain-based equity crowdfunding. They find that specific driving forces positively (e.g., costs from equity funding) and negatively (e.g., asymmetric information) affect equity funding. While Hartmann et al. [36] reveal success factors for conventional and blockchain-based crowdfunding and propose future research in this area, Stekli and Cali [69] show that equity crowdfunding via blockchain facilitates the financing of clean energy projects. Overall, blockchain technology has given entrepreneurs the capability of creating and issuing tokens for fundraising. Regulatory-compliant security tokens, including equity tokens, reduce the trust barrier that ICOs and traditional equity crowdfunding struggled with. However, equity tokens are nascent and must be designed correctly to comply with laws and regulations, ultimately reshaping the landscape of funding, entrepreneurship, and innovation [67, 80].

Method

To develop an equity token, we followed the DSR approach [32, 38, 5]. DSR generally seeks to solve an identified problem in a build-and-evaluate process that ultimately produces purposeful design artifacts [38]. Further, DSR's output can be constructs, models, methods, and instantiations, while a prototype is a typical instantiation [48]. In the end, the derived knowledge should be generalizable and, therefore, applicable to similar settings. To achieve this, we drew on both the early-stage funding and the blockchain literature when developing our blockchain prototype, deriving generalizable knowledge in a two-step evaluation. We addressed the shortfalls of the crowdfunding process and ICOs by developing and evaluating an instantiation of a blockchain-based equity token. We applied Peffers et al.'s [58] widely accepted research approach to structure our research (see Figure E.1). We iteratively used the design and development, demonstration, and evaluation steps [5, 58]. The following steps guide this research: Our research is motivated by a lack of knowledge on the design of equity tokens and their applicability. We identified traditional early-stage funding as a practically relevant problem that blockchain technology could improve [23, 31, 22, 81, 8]. We analyzed traditional equity crowdfunding problem areas and the first wave of blockchain-based solutions, i.e., ICOs. Major problems in the traditional equity crowdfunding domain include the credibility of crowdfunding platforms, a lack of secondary market trading, and high administration costs [81, 8, 63]. In contrast, ICOs pose great challenges, including missing underlying value, the need to comply with current regulations, and allowing for higher interventions [15, 79]. To address

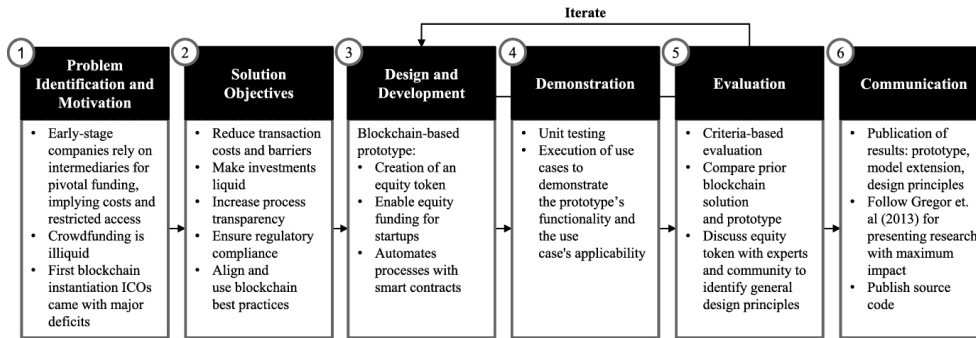


Figure E.1: Research process (adapted from Peffers et al. [58])

the identified challenges, we use both the areas for improvement of equity crowdfunding (EC-Aff) and ICOs (ICO-Aff) to derive design objectives (DOs) that an improved solution must fulfill. Furthermore, we built our derivation of DOs on the literature on equity crowdfunding and blockchain technology and the examination of past ICOs. Accordingly, we elaborate on 14 DOs for the software prototype design, implementation, and evaluation. The DOs were a starting point for the development stage. As is standard in software development, we defined the required data types and the intended solution methods. Based on the defined DOs, we implemented our equity token with additional emission and transaction protocols. We developed the prototype in an Ethereum environment since it is considered a matured platform for smart contract development [78] Finally, we conducted seven semi-structured expert interviews. This procedure allowed us to get feedback from experts on our reference implementation and the application of blockchain technology for equity tokens, which was fundamental to generalize from an instance solution to an abstract solution (see Figure E.2).

Problem Identification and Design Objectives

Limitations in the early-stage funding process are regarded as one major constraint for better exploitation of the economic potential of entrepreneurship [24, 20, 23]. In the background section, we point out several problems for early-stage equity funding raised in the equity crowdfunding literature. We argue that blockchain technology—a technology that enables trust among participants and automates business logic [30]—has the potential to address the raised deficits. ICOs promise to offer a blockchain-based alternative for crowdfunding but do not use the tokenization of equity.

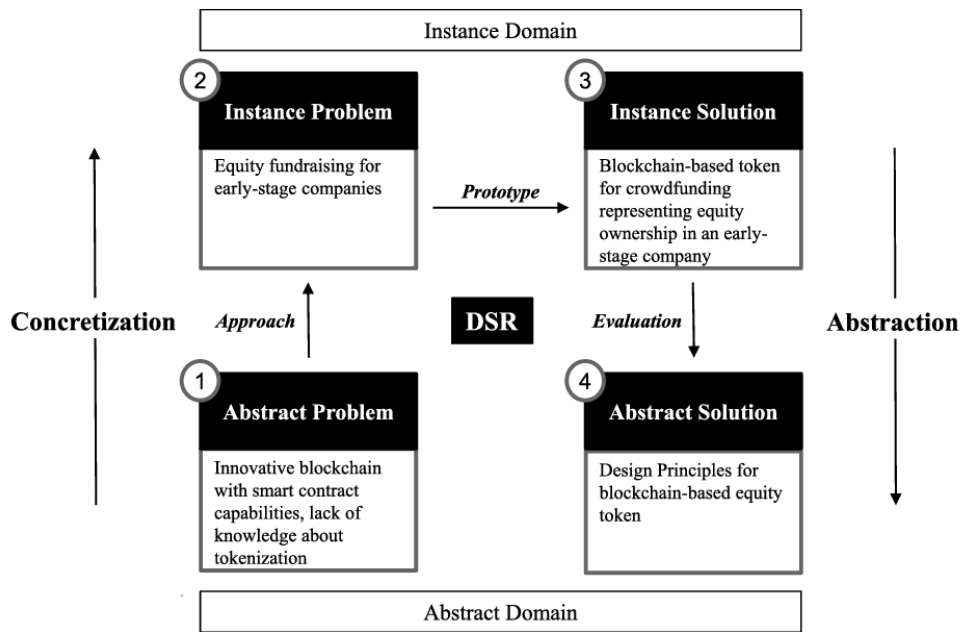


Figure E.2: Design science research: concretization and abstraction

Dimension	Area for improvement	Description of the status quo
Trust	EC-AfI01: Credibility of crowdfunding platforms	Shares of a crowdfunding company are not registered at a credible registry or traded on a reliable settlement system like public stock exchange systems. Equity crowdfunding instead relies on centralized organizations that typically have lower levels of regulation than the conventional stock market. Hence, equity crowdfunding registration and management are less secure, jeopardizing investment capital [81]

Infra- structure	EC-Afi02: Missing secondary market	In equity-based crowdfunding, investors only have limited exit options, e.g., share buy-back schemes, trade sales, or sales on the stock market after an initial public offering. In traditional stock markets, investors can sell their assets through secondary markets to other investors [68], often lacking in equity crowdfunding [28]. Thus, investors face much higher lock-in effects, limiting effective equity circulation and ultimately discouraging potential investors [81, 63]
Costs	EC-Afi03: High admin- istration/ transaction costs	Administrative processes in crowdfunding are generally based on paper documents, e.g., for registering shareholders. As investors are usually distributed regionally, there is a strong reliance on signatures and postal mail to exchange relevant documents. Such processes are time-consuming and correspondingly expensive [8, 68]
Compliance	ICO-Afi02: Not com- pliant with current regu- lations	Utility tokens do not have standard investor protections, including the ability to track ownership and identity. However, professional investors generally demand these properties [79]. In addition, existing ICOs and their token architecture on the Ethereum platform offer no built-in mechanisms for regulatory enforcement, e.g., Know-Your-Customer (KYC), Anti-Money-Laundering (AML), or token-level restrictions [64]. Although certain ICOs implemented legitimation processes, there was a lack of built-in regulation for selling these tokens to other unverified market participants [15]
Compliance	ICO-Afi03: Incompatible with higher interventions	There is a lack of operating tokens by third parties. There are several arguments for higher interventions: lost keys, unauthorized ownership, fraud, or crime—requiring access by third parties [14]

Compliance	ICO-Afi04: No reporting standards	ICO projects are often subject to an openly accessible crowd due diligence before token issuance, helping auditors retrieve transaction data early on [80, 16]. However, there is no requirement to broadcast the company’s performance upon successful funding, e.g., through quarterly financial statements, KPIs, or ad hoc messages [79]. Therefore, a multitude of ICOs does not include reporting standards
Governance	ICO-Afi05: Conflict with the entrepreneurial funding cycle	Traditionally, early-stage funding comes with a system of checks and balances to align the interests of investors and companies. In contrast, ICOs are the only funding event, with tokens often capped to realize maximum returns [12]. Thus, this model incentivizes founders to raise too much money too early, presumably leading to a waste of resources [60]
Technology	ICO-Afi06: Inflexible architecture	ICOs lack upgradability. This inflexibility in smart contract design leads to vulnerabilities. Creating future-proof smart contracts will require the ability to easily upgrade for vulnerabilities [4]
Technology	ICO-Afi07: Limited configurability	The characteristics of an investment contract for early-stage equity are manifold, owing to the variety of business models, team compositions, and different environments. ICOs—and utility tokens—offer minimal design options
Technology	ICO-Afi08: Risks in the code	The security of blockchain-based applications depends not only on the base layer but also on the smart contract. Largely, ICOs did not follow audited token standards beyond ERC20 as they were not in place [35]. Equity tokens could implement country-specific regulatory standards, which once audited enforce transaction regulations and be shared openly

Table E.2: Areas for Improvement

As the funding mechanisms show potential for improvement, we derive several AfIs from the relevant literature (see Table E.2). To ensure a practical improvement compared to conventional funding, we enrich the shortfalls de-

rived in the literature with case-specific insights from real-world funding. In particular, one of the authors conducted a conventional funding process over twelve months as the leading manager in a startup. Please note that we follow U.S. regulations when considering compliance.

Dim- ension	Design objective	Description	Evaluation criteria	Addressed Afi
Business Logic	DO01: Define and enforce specific character- istics of equity	The prototype should in- clude the key character- istics of equity regarding the entire lifecycle, e.g., is- surance, maintenance, dis- solvent, regular communi- cation (e.g., quarterly re- porting), voting rights, and equity-specific transactions (e.g., dividends)	Implemen- tation and enforce- ment	ICO-Afi01 ICO-Afi02 ICO-Afi04 ICO-Afi05
Com- pliance	DO02: Define and enforce regulatory require- ments	The prototype must com- ply with current regulatory requirements. In partic- ular, it must implement personal identification processes (KYC/AML) and token-level restric- tions. These restrictions consist of pre-transaction checks for authorized and accredited investors and require the implementation of nonfungible tokens	Implemen- tation and enforce- ment	ICO-Afi01 ICO-Afi02
Tech- nology	DO03: Provide global ac- cess to all investor types	The prototype should tech- nically allow for small in- vestments without regional censorship or discrimina- tion since it is crucial to de- mocratize investments into startups, thereby enabling funding [24]	The pos- sible number of min. invest- ment and fulfillment	Key re- quirement EC-Afi02

Compliance	DO04: Provide a framework to hamper fraud via crowd due diligence	Investors perform time and resource-consuming due diligence to assess an early-stage company's potential value owing to the large investment size. Blockchain technology allows for fragmented investments, leading to a decrease in the average investment ticket. Since this development could reduce due diligence efforts [24], the solution must allow for a uniform structured discussion and review of a venture for participants	Implementation and enforcement	EC-Afi01 ICO-Afi04
Technology	DO05: Store and process data transparently, immutably, and permanently	Data should be as transparent as possible to improve audits. Nonetheless, regulatory and data privacy considerations set reasonable boundaries for transparency. Further, to avoid malicious changes in related data, the prototype must also process transactions in a tamper-proof way and store data persistently and immutably	Fulfillment, transparency, and trust mechanisms	Key requirement EC-Afi01 ICO-Afi02 ICO-Afi04
Technology	DO06: Reduce manual activities	The manual activities involved during the issuance and management of the equity lifecycle should be automated to reduce costs and the possibilities of fraud	Manual activities	Key requirement EC-Afi03
Technology	DO07: Sufficient transactions	In an improved solution, the number of transactions should be manageable without constraints	Throughput rates	Key requirement EC-Afi02

Governance	DO08: Allow process interactions	Blockchain technology entails the ability to cut intermediaries while still incorporating governance. Thus, the solution must allow for decentralized interactions	Fulfillment	Key requirement EC-Afi03
Technology	DO09:Ensure liquidity through interoperability	The entire infrastructure stack must be fully functional before tokens can be issued and traded on secondary crypto exchanges. The development should comply with the ERC20 standard and state-of-the-art EIPs to ensure interoperability	Interoperability and tradeability	Key requirement EC-Afi02 ICO-Afi01
Governance	DO10: Align interests by implementing supervision	The prototype must issue funds with a system of checks and balances as an instrument to align the interests of different parties, e.g., voting mechanisms for board-like decisions or vesting periods	Fulfillment	EC-Afi01 ICO-Afi04 ICO-Afi05 ICO-Afi07
Business Logic	DO11: Allow multistage funding	Subsequent tranches and multistage funding allow the alignment of interests via conditional investments and further erase full funding by integrating flexible minting schemes. For multistage funding, the prototype must additionally incorporate common recapitalization mechanisms such as pro-rata	Fulfillment	ICO-Afi01 ICO-Afi05 ICO-Afi07

Tech- nology	DO12: Design the solu- tion with sufficient flexibility	To allow for the individ- ualization of investment contracts in early-stage companies, customization of the contract during the issuance process must be possible. In addition, upgradability features for unforeseeable prob- lems avoid later-stage vulnerabilities	Usage of stan- dards and proposals	ICO-Afi06 ICO-Afi08
Com- pliance	DO13: Al- low escrow and higher interven- tions	The prototype should al- low for third party opera- tions of accounts and thus specific methods to autho- rize third parties. While higher intervention is nec- essary for securities, escrow is a common steering mech- anism for special business arrangements	Fulfillment	EC-Afi01 ICO-Afi02 ICO-Afi03
Tech- nology	DO14: Embrace standards and sim- plicity	Simplicity is vital to reduce the risk of bugs and to facilitate the possibility of future adjustments. Thus, the prototype should use reviewed open-source stan- dards (e.g., ERCs, white papers) and proposals. This enhances interoper- ability, community interaction and reduces the risk of untested code on the application layer	Use of stan- dards and proposals	EC-Afi01 ICO-Afi06 ICO-Afi08

Table E.3: The Design Objectives

Based on the identified AfIs of equity crowdfunding and ICOs, we followed an iterative cycle of deriving DOs. Thus, a DO addresses one or multiple issues (AfIs) raised in the application domain. We discussed possible DOs internally and with other researchers and finally aggregated 14 DOs for our approach, which directly informs the prototype development like software requirements. For each DO evaluation, we defined criteria to evaluate the goodness of the prototype, an essential requirement for rigorous DSR research (see Table E.3).

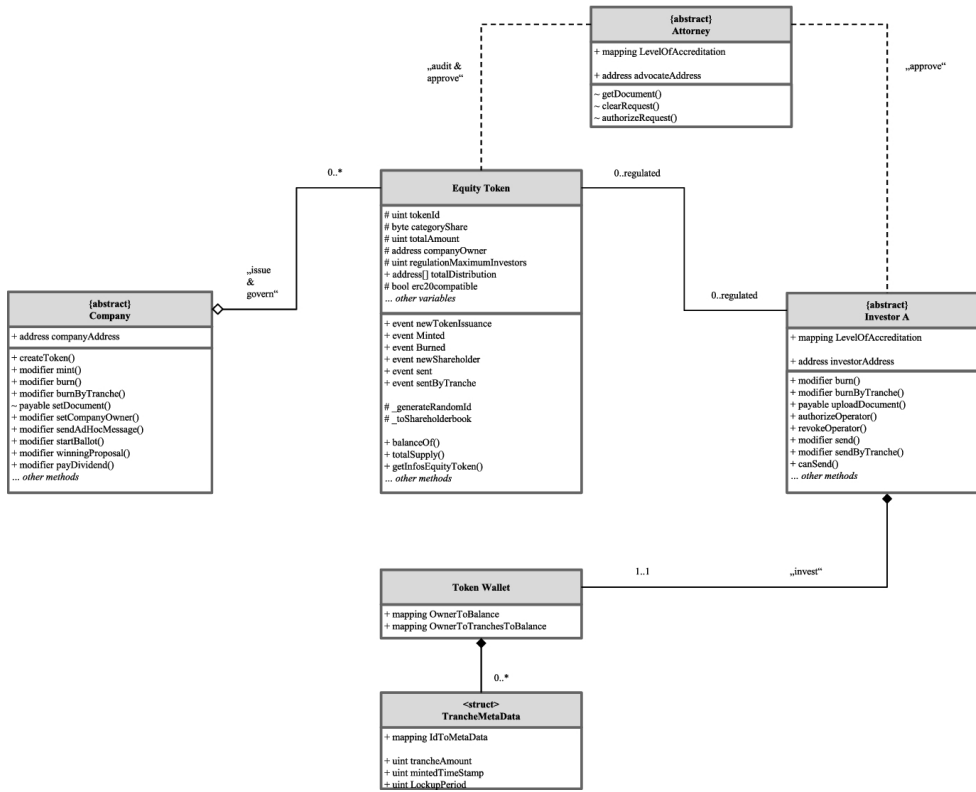


Figure E.3: Class diagram of core building blocks

Design and Development

Prototype Design and Architecture

We implemented the prototype utilizing the public and permissionless Ethereum blockchain [9]. A set of Ethereum smart contracts represent the necessary business logic. Further, we used the InterPlanetary File System (IPFS) as distributed storage technology [41] to facilitate effective document-sharing (necessary for KYC/AML). Figure E.3 illustrates the building blocks of the blockchain-based equity crowdfunding service ecosystem as a class diagram. We emphasize a core token smart contract, handling critical functionality such as transactions and accounting. Additionally, we deployed app-like smart contracts addressing the needs of different agents (issuing company, attorney, investor), such as Know-Your-Customer (KYC) and equity prospectus. As seen in Figure E.3, the core equity token implements basic functions for transferring tokens, obtaining account balances, getting the total supply of tokens, and allowing approvals. Notably, the token standard informs the core token, including the authorization layer of specific actions (modifier) [26] Each token

is implemented and deployed in a separate smart contract. This practice is common in smart contract design [25] and has several implications. Primarily, it ensures the security aspect that each funding is independent of another—a loss of access to one smart contract would not affect another. Companies can issue multiple token types over time, each with different characteristics for investors (e.g., class A or B shares), thereby addressing different investor groups. The token type is traceable by a unique identification number and is defined by pivotal metadata such as `totalAmount` and `categoryShare`, or to which `companyOwner` the token belongs. Further, a company can increase or reduce the number of previously issued tokens by issuing or burning them. The architecture ensures backward compatibility with Ethereum token standards, such as ERC20, and relatively new proposals such as EIP1400 and EIP1410 [25]. This compatibility is essential if one is to interoperate with other implementations on Ethereum. The backward compatibility can be turned on and off if new standards emerge. Once the smart contracts are deployed on the Ethereum blockchain, it assigns addresses that make the smart contracts publicly accessible. Multiple parties can then use the prototype. Only the contracts’ addresses and knowledge of the public core functionalities are required to interact with the prototype. Section 6.2. describes the token issuance and token transaction in detail. To ease the interaction with the equity token we deployed app-like smart contracts for each party. For example, the issuing company can provide necessary documents supporting the equity issuance. The documents are uploaded on IPFS and linked to a transaction on the blockchain. Furthermore, the investor can provide documents identifying himself (KYC), a necessary process that we will elaborate on in the next section.

Development and Prototype Features

Guided by Peffers et al.’s [58] DSR process and the software requirements (DOs), we developed the prototype in iterative steps following a build-and-evaluate process. For the sake of simplicity, within this paper, we demonstrate three relevant prototype features: the KYC process, the issuing process, and the transaction protocol. We selected these three features as the KYC process is a distinctive feature for equity crowdfunding in contrast to ICOs, and the token issuing process is relevant for crowdfunding in general. The last in-depth feature, token transactions, are a technical core element for transferring value on the blockchain and are of increasing importance owing to the transaction restrictions required for equity tokens. All further functionality is described in the appendix as well as documented in the open code repository.

The know-your-customer process The KYC process gains center stage for equity crowdfunding: token ownership must be continually tracked in many jurisdictions, and all investors must disclose their identities. Traditionally,

passing a KYC process conducted by a third party such as a bank or an exchange requires a potential investor for identification and final authorization. The KYC principle is crucial to fighting money laundering. Implementing the process requires that investors upload certain documents (e.g., identification documents, proof of residency), which the third party consequently authorizes. To store uploaded documents, we used IPFS, which offers the benefits of blockchain technology and is an efficient way to record documents permanently, securely, and transparently. Uploading encrypted documents with IPFS returns a hash and a key. The investor uploads the document's hash and authorizes a third party provider. Together with this message, they must send a certain fee to pay for the KYC service. The third party provider—in our example, an attorney—retrieves the documents, audits them off-chain, and either `authorizeRequest` or rejects the request. In both cases, the `accreditationFee` is automatically transferred to the third party. After approval, the investor's status code changes to `authorized`. The protocol consistently ensures that the documents can only be retrieved and encrypted by the authorized attorney. Through IPFS the investors' documents are immutably linked to the blockchain and can be tracked with the investors' address.

Token issuance At the outset, the issuer creates a token shell that determines key characteristics of the equity token, such as `tokenTicker`, `categoryShare`, and `defaultOperator`. The shell is a template for a customized equity token. Initially, the token's `totalAmount` is zero since the shell is pending, waiting for approval from a third party. For the emission of the token, documents (e.g., annual statement, prospectus) must be uploaded and audited. Again, the request passes a payable on-chain off-chain process similar to the KYC procedure. However, the required documents and audits by the attorney differ and are far more extensive. The attorney audits the shell and classifies the equity. Upon approval, the company can mint multiple rounds of this specific token, depending on its strategy, business model, and investors (see Figure E.4).

Token transaction Finally, we illustrate a transaction in detail. The transaction protocol is a key feature since equity tokens incorporate several token-level restrictions that ensure compliance with predefined regulations during the entire transaction. Thus, this design prevents accounts from transferring security tokens to unauthorized parties. Figure E.5 demonstrates the sequence diagram for a successful transaction. The issuer allocates the tokens in a primary distribution directly to the investor. Every batch of tokens in the wallet collected and controlled by an owner belongs to a unique tranche. The attached metadata describe each tranche and store information for token-level restrictions, such as a lockup period. For sending tokens, the sender can include a specific tranche for the payment or a first-in-first-out logic auto-

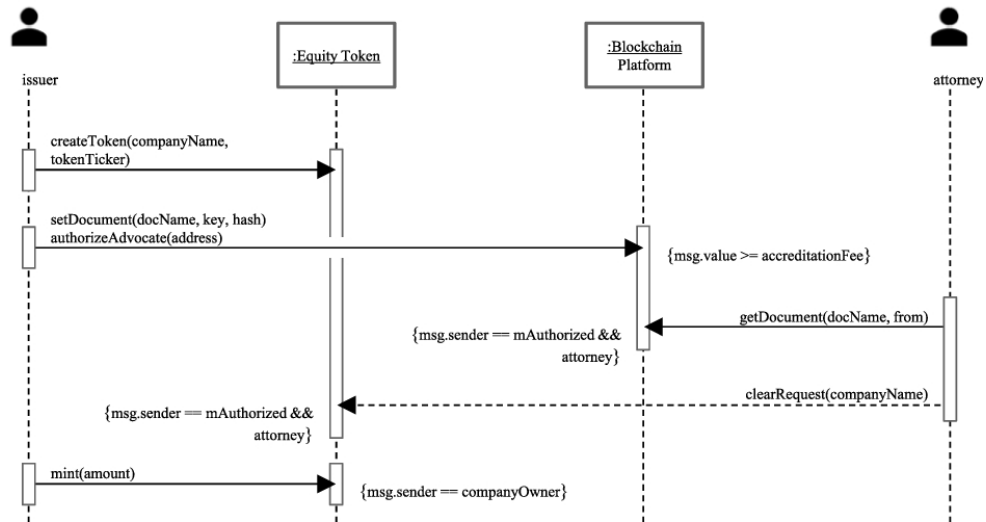


Figure E.4: Sequence diagram for the token issuance

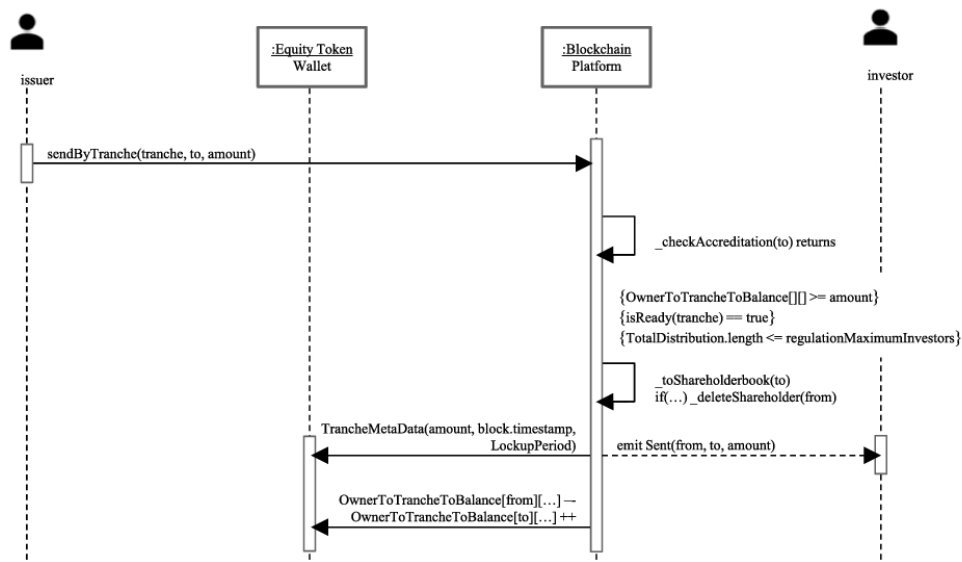


Figure E.5: Sequence diagram for a token transaction

matically selects a tranche. The sender calls `sendByTranche` and includes the receiver, amount, and tranche. The protocol then checks for both authorization (KYC/AML) and accreditation (e.g., implementing US regulations, where accreditation is conditional on the receiver's wealth) of the receiver. After the first successful check, the protocol controls whether the sender's balance is equal to or larger than the sending amount. Further, the protocol accesses both the `trancheMetaData` and general information of the token. While the tranche's metadata is necessary to check whether the lockup period has expired since the last trade, further general information allows one to check for regulatory restrictions. In our prototype, we restricted the maximum number of investors per company. The transaction protocol enlarges the public record of ownership and deletes an owner if their stake in the company is zero after a successful transaction. The receiver's wallet receives the token if all checks pass and calculates the new balance of both the sender and receiver. In the receiver's wallet, the tokens build a new tranche that gets a new specific `trancheMetaData`. As a final step, the blockchain broadcasts a successful transaction event to the network. All transactions are atomic. If only one check is unsuccessful, the blockchain will perform a rollback to the original state. A transaction can also generally be executed by an authorized operator. Only the token owner can `authorizeOperator` and `revokeOperator`, which functions as trustees to manage a portfolio. By default, a governmental address is also an authorized operator. The possibility of intervention is one mechanism to prevent fraud or crime and is a key design objective.

Evaluation

Following DSR, thoughtful evaluation of the proposed design artifact is a key request [38]. The proposed design artifact should demonstrate utility, quality, and efficacy. That is, the artifact solved the intended purpose [21]. The prototype has been fully implemented and deployed on an Ethereum test network, satisfying the core utilities in a testing environment. We proceed with a comprehensive evaluation in two steps to gauge the efficacy; each step broadened the evaluation's scope [21]. As Gregor and Hevner [32] proposed, we foremost strove for a comparative assessment, analyzing whether equity tokens are beneficial compared to the previous blockchain solution (efficacy). Thus, we applied a criteria-based evaluation and compared the prototype to the addressed Afls. Finally, we presented our research approach and prototype to industry experts in seven semi-structured interviews to evaluate the quality and derive more general insights.

Criteria-based Evaluation

We presented our prototype to the derived Afls and assessed whether the implementation of our DOs showed that an equity token improved the existing

solution (see Table E.4).

Area for improvement	Targeted DOs	Evaluation of prototype
EC-Aff01: Credibility of crowdfunding platforms	DO04 DO05 DO13 DO14	Currently, the services of platform providers conduct equity crowdfunding and thus require a significant level of trust. Given the trustless nature of Blockchain [17], participating parties can use equity crowdfunding without relying on a central intermediary
EC-Aff02: Missing secondary market	DO03 DO07 DO09	In contrast to many traditional equity crowdfunding platforms, the developed equity token enables trading on secondary markets. The prototype is based on the ERC20 interface standard and supports functions for transferring tokens to other market participants [7, 54]. The interoperability of the equity token with existing crypto exchanges and other services and existing wallets is easily realizable [54]
EC-Aff03: High administration/transaction costs	DO06 DO89	The artifact supports the digitization of the equity crowdfunding process using equity tokens as digital financial contracts. Building on IPFS allows the distribution of digital documents among the required stakeholders efficiently [41]. Therefore, especially KYC and AML-related processes could become more streamlined [47]
ICO-Aff01: Price discovery	DO01 DO02 DO09 DO11	Currently, market mechanisms drive the volatility of tokens rather than specific token designs [12]. However, the equity should become easier to price with increased market maturity, as traditional valuation models are applicable. The prototype considers the key characteristics of equity, supports the management of investor relationships during the equity's lifecycle, and ensures regulation enforcement

<p>ICO-Afi02: Not compliant with current regulations</p>	<p>DO01 DO02 DO05 DO13</p>	<p>Equity tokens generally improve investor protection since the issuance comes with legal obligations for the issuing companies. Through the programmability of equity tokens, regulatory requirements can be efficiently enforced ex-ante the transaction, saving the need to audit afterward. The prototype incorporates token-level restrictions (KYC/AML, accreditation) and provides a seamless process for investor identification. The ownership of every equity token can be continually tracked owing to partial fungibility. The prototypes offer various mechanisms to align investors' and funding seekers' interests, such as voting, vesting, multistage funding, and escrow services</p>
<p>ICO-Afi03: Incompatible with higher interventions</p>	<p>DO13</p>	<p>We implement the possibility of access by a third party operator into the equity token prototype. This optional mechanism enables higher interventions in instances of fraud or unauthorized ownership. Our prototype also allows for the authorizing and revoking of operators</p>
<p>ICO-Afi04: No reporting standards</p>	<p>DO01 DO04 DO05 DO12</p>	<p>Our prototype enables ad hoc messages by the issuing company but does not enforce them regularly with oracles. Nonetheless, we argue that it should be discussed whether communication must be stored on-chain or can be managed off-chain</p>
<p>ICO-Afi05: Conflict with the equity funding cycle</p>	<p>DO12 DO13 DO13</p>	<p>Through increased rights and obligations, equity tokens provide a higher level of alignment. Equity tokens can be used to financially incentivize through tranches, different equity token types, as well as multiple batches of one type, can be issued conditionally on specific milestones</p>
<p>ICO-Afi06: Inflexible architecture</p>	<p>DO12 DO14</p>	<p>The prototype implements a proxy logic to swap smart contracts in case of updates. However, this is not optimal as smart contracts' upgradability on a blockchain is a general problem [74], providing future research opportunities</p>
<p>ICO-Afi07: Limited configurability</p>	<p>DO10 DO11</p>	<p>Through programmability, equity tokens offer an infinite design space for securities' financial engineering and technical implementation [17]</p>

ICO-Aff08: Risks in the ICO code	DO12 DO14	Most utility tokens follow the same token standard (i.e., ERC20). This standard can also be applied to equity tokens. On the other hand, the increased complexity of equity tokens increases the risk of malicious code. Our prototype incorporates state-of-the-art proposals (e.g., ERC20, ERC777, EIP1400) and thus also helps to standardize equity tokens
----------------------------------------	--------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Table E.4: Criteria-based Evaluation of Affs and DOs

In sum, many DOs seek to enhance trust and reduce adverse selection impacts by dismantling the asymmetrical information between interacting parties, aligning interests, and minimizing the regulatory uncertainty about an equity token. Equity tokens reduce the overall transaction costs of early-stage funding. Decentralization is a fundamental benefit of blockchain, reducing the middlemen and expenses required to conduct transactions on the Ethereum blockchain.², at more than 20 transactions per second [6]. In general, we propose that token funding changes the market’s perspective: traditionally, funding-seekers must discuss funding terms with every single potential investor. Using equity tokens improves efficiency since the issuers’ terms are broadcasted worldwide via the blockchain and accompanied by real-time settlement. Overall, the implemented DOs reduced transaction costs for purchasing and trading in equity and technically granted access to investors type globally. Furthermore, small investments become economically viable owing to lower transaction costs. Token-level restrictions and investor identification ensure high compliance levels and thus secure the underlying value of a security on the blockchains. Overall, the transparency increases since each update of the equity token’s implementation include a timestamp recorded on the blockchain and stores key documents publicly.

Semi-structured Interviews

We conducted seven semi-structured interviews with industry experts to evaluate our prototype for quality and derive generalized design principles for equity tokens. For our research approach, semi-structured interviews are a natural fit since they are a flexible evaluation technique. On the one hand, the interviewer sets up a general interview structure and covers the main questions, deciding in advance on the direction to be covered; on the other hand, the interviewee has a fair degree of freedom on how to answer and to what extent [19, 56] We reached out to potential interview partners from the authors’ network. In general, we aimed to gather a heterogeneous interview

²The test scenarios yielded average computational costs of 821,000 gas for creating and minting tokens (without one-time KYC and AML)

Table E.5 Overview of the interviewees

Expert ID	Professional title	Field of expertise	Organization type	Relevant experience
EXP1	COO & Entrepreneur	Blockchain early-stage funding	Research institute	>8y
EXP2	Fund Manager	Early-stage funding, crowdfunding	Investment bank	>3y
EXP3	COO & Consultant	IT platform, Blockchain	Crypto exchange	>5y
EXP4	Research Assistant	Blockchain	Research institute	>3y
EXP5	Senior Consultant	Technology transformation	IT Consultancy	>5y
EXP6	Business Developer	Blockchain	Blockchain community	>3y
EXP7	Head of Sales	Blockchain, crowdfunding	Blockchain fintech	>8y

panel, including academics, practitioners, and technical or business experts. In total, we conducted two rounds of interviews: starting with three interviewees and adding four more experts in the second iteration (reducing the interviewees' time commitment). The interviews took place at the end of 2020, and the participants are listed in Table E.5. Beforehand, all the interviewees received a summary presentation about the research approach, the underlying problem domain, and crucial working definitions to foster open discussion and maximize the output. In the structured part of the interviews, we discussed the lists of Afls and DOs. The interviewees assessed the Afls and DOs according to agreement, performance, prioritization, and completeness. We included the results of this feedback directly into our design artifact, utilizing the iterative nature of our research approach (see Figure E.6), which has proven beneficial multiple times in IS research [5, 58]. The semi-structured part of the interview consisted of a set of open questions to allow for open discussion of all aspects. The twelve questions have been created in multiple workshops among the author team. Questions included the advantages and disadvantages of blockchain-based tokens for equity crowdfunding, the value-add of blockchain technology within the crowdfunding process, and the technology's maturity and biggest remaining hurdles. We sought to achieve a more general understanding of blockchain-based equity crowdfunding, facilitating a higher abstraction level and deriving more general applicable knowledge. We recorded the interviews and used qualitative techniques, such as the transcription and coding of the interviews. Later, the authors discussed

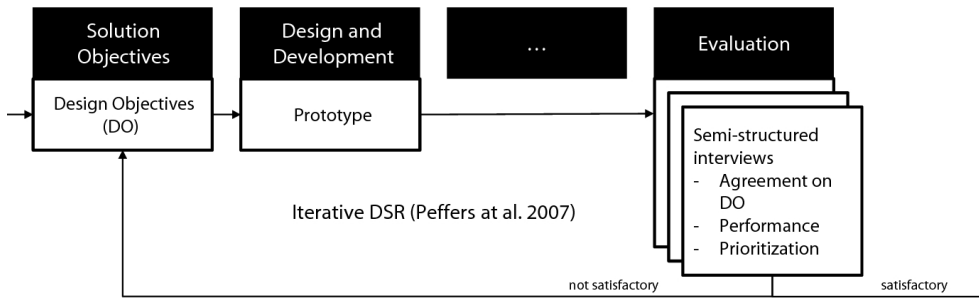


Figure E.6: Iterative design, development, and evaluation of the artifact (based on Peffer et al. [58])

the results of the analysis until a common understanding was reached. All interviewees emphasized that blockchain technology can play a crucial role in early-stage equity funding if the funding seekers' applicability becomes more convenient and fully exploits blockchain technology's benefits. In addition, the interviewees agreed that the following key attributes exploit tokenization's potential fully: increased liquidity, divisibility, reduced friction, disintermediation, removed geographical barriers, and more transparency. Interestingly, every expert acknowledged that the Ethereum blockchain provides a matured infrastructure for developing equity tokens. EXP7 stated that this is particularly true since Ethereum enables the implementation of smart contracts, has a larger development community, features more robust IT security, and allows for the compatibility of token standards. Concerning privacy, EXP4 agreed to use Ethereum and recommended considering a permissioned blockchain such as Hyperledger Fabric since it provides built-in privacy features. To address the prototype's applicability, they called for reducing the technical entry barriers of equity tokens through a customer-friendly user experience and further standardization of protocols. The interviewees mentioned unclear and fragmented regulations as one primary challenge to exploiting the full potential of equity tokens calling for a clean regulatory environment without limiting innovation in this space. Concerning the transformation from ICOs to equity tokens, all are seeing a considerable improvement compared to the first wave of blockchain funding and agreed to strict definitions determined by the token characteristics. All the interviewees valued improved investor protection, token-level regulations, and the underlying value of security tokens. In this context, EXP2 stated that volatility and speculation owing to immature valuation was also a phenomenon in equity during the Dot-Com bubble. But with a maturing market, the valuation methods and experience improved. Regarding ICOs, EXP3 stated that the financial success was faster than the technology's maturity and emphasized that ICOs addressed "retail investors without time or an interest in doing due diligence." While the pub-

lic has pushed ICOs, he expects that the established industry's equity tokens will be valued more rationally. Indeed, EXP3 called it a "desirable development" since "retail investors should not be in that space." Market liquidity for equity tokens was another key discussion with all the interviewees. EXP7 supported stated that tokenization is especially useful when considering asset classes with low trading volumes as large assets are already trading efficiently. In this context, EXP6 said that tokenization "makes dead capital" (i.e., illiquid asset classes, such as crowdfunding) more liquid, and allows for fractional ownership, ultimately granting access to a broader investor base. Also, the regulation of equity tokens was a controversial topic among the interviewees. While they all agreed that a certain level of regulation is necessary for equity tokens, the optimal level of regulation they proposed was diverse. EXP4 noted that, in this context, it is crucial to grant access to various participants, such as tax authorities, brokers, exchanges, and other financial services, and to set standards that are supported by public authorities. Such an approach could also include the use of master keys, allowing for the freezing of assets. EXP6, on the other hand, denied the meaningfulness of allowing central entities to take corrective actions: "this would counteract the whole idea of blockchain, making a decentralized system central again." EXP5 eventually pointed out that regulating equity tokens is a mixed bag. While handling AML requires master keys, over-regulation can lead to tokens losing their benefits compared to conventional systems. Following EXP5, technical standards are strictly required to allow for the mass adoption of equity tokens. Remarkably, the ERC20 demonstrates the effect of agreeing on a specific standard, facilitating a substantial number of ICOs. Further, standards are necessary to integrate third parties, such as exchanges. This interviewee emphasized the nascent status quo and called for further development in this field.

Discussion

Our design has introduced an approach for automated, secure, and customized issuance of an equity token on the Ethereum blockchain, aiming to provide a novel approach for equity crowdfunding. Thus, we contribute to the body of knowledge on the developing blockchain-based equity crowdfunding domain [37, 81, 36, 1, 69].

The literature on equity crowdfunding points out that investors only have limited exit options, leading to higher risks and despair [81, 68, 28]. Conversely, our system allows an early-stage company to create and distribute their shares on a primary issuance platform and facilitates interfaces to exchanges for secondary market trading. In addition, investor relationships can be managed by the issuing company on-chain throughout equity lifecycle applications. Since every successful transaction of tokens is automatically recorded, the system provides a complete and tamper-proof transaction history and distribution

of the equity token ownership. The system operates without institutional involvement through decentralized protocols and complies with a predefined regulatory framework, owing to self-regulating tokens. We find that by using an instance of a blockchain-based equity token for crowdfunding, the advantages of tokenizing equity can be realized, as demonstrated with our prototype, and therefore agree with Chen [13] and Roth et al. [63].

From a technological perspective, applying blockchain technology in the equity domain constitutes multiple benefits. First, due to decentralized protocols, trusted institutional intermediaries are not necessary to manage the system infrastructure like accounts and transactions, thereby largely reducing friction [45]. Not a single participant in the system needs to be trusted because the inherent consensus mechanism of blockchains ensures the network's administration and follows smart contracts' logic. Inadequate use is still possible but is lowered to a minimum since the deployed algorithms govern human behavior [6, 81]. Second, blockchains' decentralized structure allows us to store all the relevant data on the network's nodes [16]. Thus, our prototype enhanced general reporting and auditability since the nodes store all relevant data transparently and allow regulatory entities or third party providers to retrieve them easily. Due to its high level of redundancy, the system becomes resilient against potential cyberattacks and prevents single points of failure [6]. The inherent security features of our equity token reduce the trust barrier in crowdfunding, which remained a major concern in traditional, centralized equity crowdfunding [81]. Overall, our system works like a transparency device that assures the availability of a complete, valid, and public record of both historical and present equity ownership, thereby encrypting and attaching key documents (e.g., KYC) [53, 61]. Third, the prototype significantly reduces transaction processing time since blockchain uniquely combines the recording and value transaction. Traditional equity crowdfunding suffers from cumbersome administration processes. These include paper-based documentation and global distribution by mail, which dramatically slows down the transfer of ownership and thus increases dependencies on intermediaries [68]. Our equity token clears initial transactions in seconds, thus fostering the rapid exchange of ownership.

Overall, our understanding of blockchain-based equity crowdfunding differs from traditional equity crowdfunding. Thus, we extend the traditional crowdfunding model proposed by Haas et al. [33] by redesigning the service ecosystem holistically through the introduction of blockchain in the context of equity crowdfunding. Our model extension reveals the elimination of payment providers and banks through blockchain, which now covers all services provided by the former intermediaries (see Figure E.7). To consider the regulatory requirements of equity crowdfunding, we also include attorneys, regulatory authorities, and external auditors as vital stakeholders within the system. We correspondingly note that our model also differs from the one proposed by Schweizer et al. [67]. While there are differences, as Schweizer et al. [67]

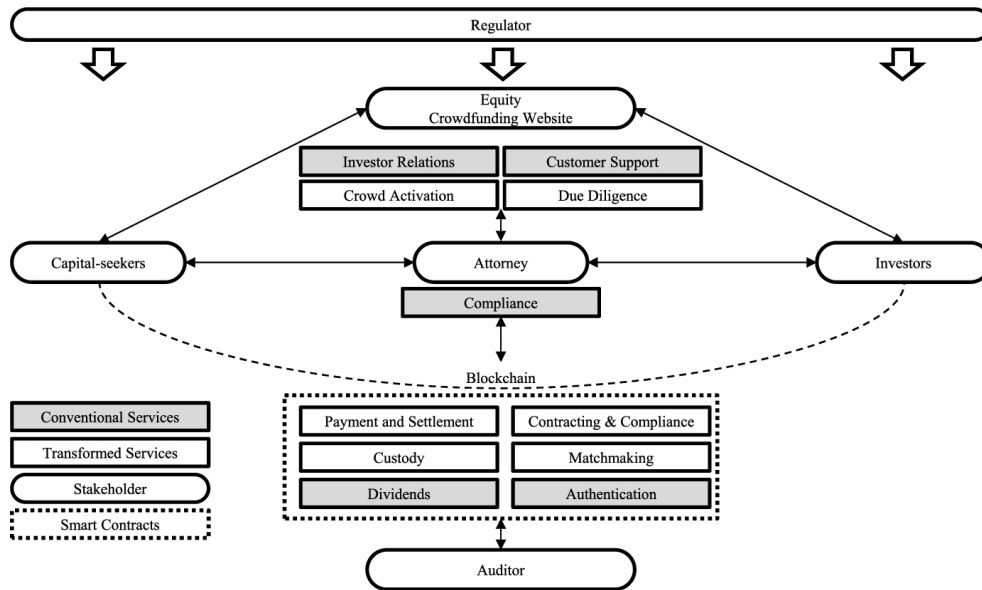


Figure E.7: Blockchain equity crowdfunding service ecosystem (based on Haas et al. [33])

describe their model in the context of crowdlending, we disagree with the general conception that blockchain entirely cuts out intermediaries and all tasks can be outsourced. In essence, their model shows that smart contracts can be responsible for all services provided by the crowdfunding partner, including crowd activation and customer support. Although it is theoretically possible to outsource these tasks to very complex smart contracts, we still see the crowdfunding partner as an essential stakeholder to provide the mentioned services. Similar to exchanges providing services on top of ICO tokens, we propose that crowdfunding partners offer services on top of equity tokens, e.g., due diligence. Besides, operations that financial institutions previously managed, such as authentication, custodial services, and dividend payouts, are now automated through smart contracts. Furthermore, blockchain facilitates instant clearance and settlement of payments, removing transaction friction. Design science should provide archival knowledge [38], and, thus, contributing to design theory is a vital part of conducting DSR [32]. Following Beck et al. [6], we propose design principles (see Table E) to contribute to the body of knowledge on designing blockchain-based systems [6, 32]. Due to extensive prototyping, rigorous evaluation, and semi-structured interviews, we generalize our findings and thus argue applying equity tokens beyond the equity crowdfunding domain. Accordingly, they could act as comprehensible guidelines for the effective design of equity tokens.

DP1: Leverage a combination of blockchain and other distributed

Table E.6 Design Principles for Blockchain-based Equity Tokens

Design Principle	Addressed Afls
DP1 Lever a combination of blockchain and other distributed technologies	EC-Afl01, EC-Afl03, ICO-Afl02, ICO-Afl04
DP2 Lever token metadata to include granular transaction requirements	ICO-Afl01, ICO-Afl02, ICO-Afl07
DP3 Follow token standards and standard interfaces to increase interoperability	EC-Afl01, EC-Afl02, EC-Afl03, ICO-Afl02, ICO-Afl04, ICO-Afl08,
DP4 Central administration should only be incorporated as a last resort	EC-Afl01, ICO-Afl02, ICO-Afl03, ICO-Afl08
DP5 Allow for multiple tranches over the token life cycle	EC-Afl02, ICO-Afl01, ICO-Afl02, ICO-Afl05, ICO-Afl07
DP6 Use a public blockchain to facilitate transparency	EC-Afl01, ICO-Afl04, ICO-Afl08,
DP7 Give power to the machine	EC-Afl01, ICO-Afl03, ICO-Afl04, ICO-Afl06, ICO-Afl08

technologies Off-chain physical documents are often needed to assess a claim of ownership. Notably, progressive jurisdictions are moving forward to replace physical documents with digital ones. To minimize the data necessary to be stored on a blockchain (and thus costs), we advise storing a pointer (i.e., hash) toward a set of documents instead of storing the documents. In particular, distributed systems such as IPFS can build a suitable balance between the complete centralization of legacy systems and highly decentralized public blockchains.

DP2: Lever token metadata to include granular transaction requirements Every equity token should include metadata. Thus, equity tokens can become fungible. Metadata is a pre-requirement to set up very granular transaction conditions, which can be asserted with every transaction. Incorporating transaction requirements in smart contracts allows checking requirements before a transaction is executed. This assertion renders post-transaction audits completely obsolete. Typical examples are the accreditation status, the token creation date, or the emitting jurisdiction.

DP3: Follow token standards and standard interfaces to increase interoperability The blockchain, a single infrastructure layer, powers crypto tokens. In applying the same standards to the token, these assets can interact with one another. Standards can be established by open-source communities, corporate alliances, and academia or can be determined by governments. For instance, in our Ethereum prototype, ERC20 (token standard) and EIP1400

(security token standard) received significant community support. Interoperability eventually increases the entire ecosystem’s efficiency. Additionally, open standards reduce the chances of security flaws through peer code reviews.

DP4: Central administration should only be incorporated as a last resort The reason for central administration is manifold. Regulation (e.g., AML) and security flaws (e.g., as it happened with the DAO hack) require centralized entities to intervene. As such, we implemented options to register public keys, which allows the owner to pause tokens. While we acknowledge the necessity for such centralized administration, we still consider it a last resort method since it directly goes against a vital feature of a blockchain—decentralization.

DP5: Allow for multiple tranches over the token life cycle The practice of attaching metadata to equity tokens and technically structuring the tokens according to their metadata, i.e., tranching, is beneficial to allow very granular token transaction requirements (see DP2) and supports the issuance of differently designed equity tokens over the lifecycle of a company. As early-stage companies are dynamic and have multiple funding rounds, each round could be represented by a new tranche of equity tokens.

DP6: Use a public blockchain to facilitate transparency Public blockchain technology is inherently transparent as it stores transactions publicly and immutably on a distributed register. By design, this transparency results in the public recording of all equity token transactions. The companies’ equity management, such as dividend payments or issuance of new tokens, is stored throughout the lifecycle. This implementation potentially decreases the burden on reporting and auditing of the company.

DP7: Give power to the machine Smart contracts allow the automation of arbitrary business logic securely. Therefore, we promote their use to automate recurring tasks of equity tokens. For example, in the prototype, we used sophisticated transaction restriction assertions: It is technically infeasible to send the equity token to a non-compliant receiver. Outsourcing automation to smart contracts potentially increases efficiency as well as system robustness. Thorough one-time audits ensure that smart contracts are always executed correctly.

We position our research to fill the gap in the IS literature on the design theory of blockchain-based equity tokens. We used a rigorous DSR approach to the design, development, and evaluation of a blockchain-based equity token prototype for crowdfunding [32, 58]. Thus, we answer our research questions on how blockchain can be incorporated as an alternative infrastructure for equity crowdfunding. In addition, we extended an established crowdfunding model and developed seven principles for the effective design of equity tokens. Overall, we embedded our theoretical insights in the current academic discourse, thereby following the calls by Treiblmaier et al. [76], Kranz et al. [42], and Perdana et al. [59] to contribute to the design theory on blockchain tokens.

Conclusion

The developed blockchain prototype sought to offer new insights into the design of equity tokens. We designed an instance solution for the problem areas of equity crowdfunding and ICOs, developing an equity token that covers the entire equity lifecycle. We derived general knowledge that is eventually applicable to blockchain-based equity beyond equity crowdfunding through the development, evaluation, and expert interviews. We sought to make several contributions to the body of knowledge. First, by focusing on a specific form of company funding and presenting the solution design, we provided an answer to effectively tokenizing equity for crowdfunding. Second, the research process helped us better understand whether a particular type of crowdfunding could benefit from the characteristics of blockchain. Third, we provided an extended model for the blockchain-based equity crowdfunding service ecosystem. Fourth, we derived generalized design principles to guide the design and development of blockchain-based equity tokens. In addition, our research offered various practical implications. First, early stage companies can use the source code of our prototype to build an equity token to fund their business, thus, improving the funding process holistically. Second, we showed that certain third parties will still play an essential role in the early-stage funding ecosystem regarding the complex regulatory requirements. Third, the prototype demonstrated how using the token ecosystem could increase the liquidity of equity shares and encourage secondary market trading, opening the equity market to new investors. Also, this study had limitations. We used the Ethereum blockchain as an instantiation reference. However, public blockchains could function as an infrastructure with improved privacy features and performance. Although we have provided an instantiation example, our design principles required additional validation with qualitative interviews backing our findings. We, therefore, call for future research into understanding the relationships between company funding and the benefits equity can gain from being tokenized. In addition, future design-oriented research could apply our design principles in different contexts, e.g., private equity or venture capital, and can thus assess their general applicability. Equity token will establish their places in the blockchain ecosystem considering the rapid development and increased interest in the equity token ecosystem. In the following years, we expect that many equity tokens will enter the market. From a technological perspective, equity tokens have substantial potential to improve legacy financial infrastructures vastly. From a business perspective, equity tokens will facilitate the funding process.

Paper References

- [1] Asma' Tajul Arifin, Nur Aishah Arshad, and Aishath Muneeza. The application of blockchain technology in crowdfunding: Towards financial inclusion via technology. *International Journal of Management and Applied Research*, 5(2):82–98, 2018.
- [2] Laurin Arnold, Martin Brennecke, Patrick Camus, Gilbert Fridgen, Tobias Guggenberger, Sven Radszuwill, Alexander Rieger, André Schweizer, and Nils Urbach. Blockchain and initial coin offerings: Blockchain's implications for crowdfunding. In *Business Transformation through Blockchain*, pages 233–272. Palgrave Macmillan, Cham, Switzerland, 2019.
- [3] Nina Bachmann, Benedict Drasch, M. Miksch, and A. Schweizer. Dividing the ico jungle: Extracting and evaluating design archetypes. In *Wirtschaftsinformatik*, 2019.
- [4] A. Banerjee, A. Belyaeva, C. Frankopan, M. Mersch, and R. Muirhead. The state of the token market, 2017. URL <https://static1.squarespace.com/static/5a19eca6c027d8615635f801/t/5a73697bc8302551711523ca/1517513088503/The+State+of+the+Token+Market+Final2.pdf>.
- [5] Roman Beck, Sven Weber, and Robert Wayne Gregory. Theory-generating design science research. *Information Systems Frontiers*, 15(4):637–651, 2013.
- [6] Roman Beck, Jacob Stenum Czepluch, Nikolaj Lollike, and Simon Malone. Blockchain – the gateway to trust-free cryptographic transactions. *European Conference on Information Systems (ECIS 2016)*, 2016.
- [7] Garry Bruton, Susanna Khavul, Donald Siegel, and Mike Wright. New financial alternatives in seeding entrepreneurship: Microfinance, crowdfunding, and peer-to-peer innovations. *Entrepreneurship Theory and Practice*, 39(1):9–26, 2015.
- [8] B. Buerger, Andreas Mladenow, and Christine Strauss. Equity crowdfunding market: assets and drawbacks. In *International Conference on Information Systems (ICIS 2017)*, pages 1–6, 2017.

- [9] Vitalik Buterin. Ethereum white paper: a next generation smart contract & decentralized application platform, 2014. URL <https://ethereum.org/en/whitepaper/>.
- [10] Cambridge Center for Alternative Finance. The global alternative finance market benchmarking report: Trends, opportunities and challenges for lending, equity, and non-investment alternative finance models, 2020. URL <https://www.jbs.cam.ac.uk/wp-content/uploads/2020/08/2020-04-22-ccaf-global-alternative-finance-market-benchmarking-report.pdf>.
- [11] Martin A. Carree and A. Roy Thurik. The impact of entrepreneurship on economic growth. In Zoltan J. Acs and David B. Audretsch, editors, *Handbook of Entrepreneurship Research*, pages 557–594. Springer New York, New York, NY, 2010.
- [12] Christian Catalini and Joshua S. Gans. Initial coin offerings and the value of crypto tokens. *SSRN Electronic Journal*, 2019. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3137213.
- [13] Yan Chen. Blockchain tokens and the potential democratization of entrepreneurship and innovation. *Business Horizons*, 61(4):567–575, 2018.
- [14] Jiri Chod and Evgeny Lyandres. A theory of icos: Diversification, agency, and information asymmetry. *SSRN Electronic Journal*, 2018.
- [15] U. W. Chohan. Initial coin offerings (icos): Risks, regulation, and accountability. In Stéphane Goutte, Khaled Guesmi, and Samir Saadi, editors, *Cryptofinance and Mechanisms of Exchange*, Contributions to Management Science Ser, pages 165–177. Springer International Publishing, Cham, Switzerland, 2020.
- [16] K. Christidis and M. Devetsikiotis. Blockchains and smart contracts for the internet of things. *IEEE Access*, 4:2292–2303, 2016.
- [17] Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797, 2019.
- [18] Andy Cosh, Douglas Cumming, and Alan Hughes. Outside entrepreneurial capital. *The Economic Journal*, 119(540):1494–1533, 2009.
- [19] Alfonso de la Rocha. Anatomy of an erc: An exhaustive survey - coinmonks - medium. *Coinmonks*, 7.5.2018. URL <https://medium.com/coinmonks/anatomy-of-an-erc-an-exhaustive-survey-8bc1a323b541>.
- [20] David J. Denis. Entrepreneurial finance: an overview of the issues and evidence. *Journal of Corporate Finance*, 10(2):301–326, 2004.

- [21] Benedict J. Drasch, Gilbert Fridgen, Tobias Manner-Romberg, Fenja M. Nolting, and Sven Radszuwill. The token’s secret: the two-faced financial incentive of the token economy. *Electronic Markets*, 30(3):557–567, 2020.
- [22] Steven Dresner. *Crowdfunding: A Guide to Raising Capital on the Internet (Bloomberg Financial)*. Wiley, 2014.
- [23] Claire Economidou, Luca Grilli, Magnus Henrekson, and Mark Sanders. Financial and institutional reforms for an entrepreneurial society. *Small Business Economics*, 51(2):279–291, 2018.
- [24] Saul Estrin, Daniel Gozman, and Susanna Khavul. The evolution and adoption of equity crowdfunding: entrepreneur and investor entry into a new market. *Small Business Economics*, 51(2):425–439, 2018.
- [25] Ethereum. Home — ethereum improvement proposals, 2020. URL <https://eips.ethereum.org/>.
- [26] Brad Feld and Jason Mendelson. *Venture Deals*. John Wiley & Sons, Inc, Hoboken, NJ, USA, 2016.
- [27] Gianni Fenu, Lodovica Marchesi, Michele Marchesi, and Roberto Tonelli. The ico phenomenon and its relationships with ethereum smart contract environment. In Roberto Tonelli, Stéphane Ducasse, Gianni Fenu, Andrea Bracciali, and IEEE International Workshop on Blockchain Oriented Software Engineering, editors, *2018 IEEE 1st International Workshop on Blockchain Oriented Software Engineering (IWBOSE)*, pages 26–32, [Piscataway, NJ], 2018. IEEE.
- [28] Freedman M. Freedman. The growth of equity crowdfunding, 2015. URL <https://www.primerus.com/wp-content/uploads/2015/08/Coleman-and-Horowitz-LLP-The-Growth-of-Equity-Crowdfunding.pdf>.
- [29] Gilbert Fridgen, Sven Radszuwill, Nils Urbach, and Lena Utz. Cross-organizational workflow management using blockchain technology - towards applicability, auditability, and automation. In *Proceedings of the 51st Hawaii International Conference on System Sciences*, 2018.
- [30] Florian Glaser. Pervasive decentralisation of digital infrastructures: A framework for blockchain enabled system and use case analysis. In *Proceedings of the 50th Hawaii International Conference on System Sciences (2017)*, 2017.
- [31] Paul Gompers and Josh Lerner. *The venture capital cycle*. The MIT Press, Cambridge, Massachusetts, USA, 2004.

- [32] Shirley Gregor and Alan R. Hevner. Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2):337–355, 2013.
- [33] Philipp Haas, Ivo Blohm, Christoph Peters, and Jan Marco Leimeister. Modularization of crowdfunding services: Designing disruptive innovations in the banking industry. In *International Conference on Information Systems (ICIS 2015)*, 2015.
- [34] Christopher Hahn and Adrian Wons. *Initial Coin Offering (ICO): Unternehmensfinanzierung auf Basis der Blockchain-Technologie*. Springer Gabler, Wiesbaden, 2018.
- [35] Stephen Hall. 6 steps to erc20 tokens and ico smart contracts - coinmonks - medium. *Coinmonks*, 7.8.2018. URL <https://static1.squarespace.com/static/5a19eca6c027d8615635f801/t/5a73697bc8302551711523ca/1517513088503/The+State+of+the+Token+Market+Final2.pdf>.
- [36] Felix Hartmann, Gloria Grottolo, Xiaofeng Wang, and Maria Ilaria Lunesu. Alternative fundraising: Success factors for blockchain-based vs. conventional crowdfunding. In Robert Tonelli, editor, *IWBOSE '19*, pages 38–43, Piscataway, NJ, 2019. IEEE.
- [37] Felix Heieck, Tatiana Ermakova, Benjamin Fabian, and Stefan Lessmann. Equity crowdfunding based on the blockchain? a delphi study. *SSRN Electronic Journal*, 2018. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3198083.
- [38] Alan Hevner, Salvatore March, Jinsoo Park, and Sudha Ram. Design science in information systems research. *Management Information Systems Quarterly*, 28(1), 2004.
- [39] Garrick Hileman and Michel Rauchs. Global cryptocurrency benchmarking study. *Cambridge Centre for Alternative Finance*, 33:33–113, 2017.
- [40] ICOData.de. Icodata, 2021. URL <https://www.icodata.io/>.
- [41] IPFS. Ipfs powers the distributed web, 2020. URL <https://ipfs.io/>.
- [42] Johann Kranz, Esther Nagel, and Youngjin Yoo. Blockchain token sale. *Business & Information Systems Engineering*, 61(6):745–753, 2019.
- [43] Jiasun Li and William Mann. Initial coin offering and platform building. *SSRN Electronic Journal*, 2018. URL [10.2139/ssrn.3088726](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3320884).
- [44] Dan Liebau and Patrick Schueffel. Crypto-currencies and icos: Are they scams? an empirical study. *SSRN Electronic Journal*, 2019. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3320884.

- [45] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making smart contracts smarter. In Edgar Weippl, editor, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 254–269, New York, NY, 2016. ACM.
- [46] Evgeny Lyandres, Berardino Palazzo, and Daniel Rabetti. Are tokens securities? an anatomy of initial coin offerings. *SSRN Electronic Journal*, 2018. URL https://www.researchgate.net/publication/329719373_Are_Tokens_Securities_An_Anatomy_of_Initial_Coin_Offerings.
- [47] Abdullah Al Mamun, Abdullah Al Mamun, Sheikh Riad Hasan, Md Salahuddin Bhuiyan, M. Shamim Kaiser, Mohammad Abu Yousuf, and Mohammad Abu Yousuf. Secure and transparent kyc for banking system using ipfs and blockchain technology. In *2020 IEEE Region 10 Symposium (TENSYP)*, pages 348–351. IEEE, 2020.
- [48] Salvatore T. March and Gerald F. Smith. Design and natural science research on information technology. *Decision Support Systems*, 15(4): 251–266, 1995.
- [49] Jens Mattke, Christian Maier, Lea Reis, and Tim Weitzel. Bitcoin investment: a mixed methods study of investment motivations. *European Journal of Information Systems*, pages 1–25, 2020.
- [50] Florie Mazzorana-Kremer. Blockchain-based equity and stos: Towards a liquid market for sme financing? *Theoretical Economics Letters*, 09(05): 1534–1552, 2019.
- [51] Alexandra Moritz and Joern Hendrich Block. Crowdfunding und crowd-investing: State-of-the-art der wissenschaftlichen literatur (crowdfunding and crowdinvesting: A review of the literature). *SSRN Electronic Journal*, 2013. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2274141.
- [52] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. URL <http://www.bitcoin.org/bitcoin.pdf>.
- [53] Benedikt Notheisen, Jacob Benjamin Cholewa, and Arun Prasad Shanmugam. Trading real-world assets on blockchain. *Business & Information Systems Engineering*, 59(6):425–440, 2017.
- [54] Luis Oliveira, Liudmila Zavolokina, Ingrid Bauer, and Gerhard Schwabe. To token or not to token: Tools for understanding blockchain tokens. In *International Conference on Information Systems (ICIS 2018)*, 2018.
- [55] Francisco-Javier Olleros and Majlinda Zhegu, editors. *Research Handbook on Digital Transformations*. Edward Elgar Publishing, Northampton, Massachusetts, USA, 2016.

- [56] OpenZeppelin. Contracts - openzeppelin docs, 2020. URL <https://docs.openzeppelin.com/contracts/3.x/>.
- [57] Aleksei Panin, Kai-Kristian Kemell, and Veikko Hara. Initial coin offering (ico) as a fundraising strategy: A multiple case study on success factors. In *Software Business*, volume 370 of *Lecture Notes in Business Information Processing*, pages 237–251. Springer International Publishing, Cham, Switzerland, 2019.
- [58] Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3):45–77, 2007.
- [59] Arif Perdana, Alastair Robb, Vivek Balachandran, and Fiona Rohde. Distributed ledger technology: Its evolutionary path and the road ahead. *Information & Management*, 2020.
- [60] Denis Petrovic. Icos, don’t bite off more than you can chew — hacker noon, 2017. URL <https://hackernoon.com/icos-dont-bite-off-more-than-you-can-chew-d658aae9579e>.
- [61] Omri Ross, Johannes Jensen, and Truls Asheim. Assets under tokenization. In *International Conference on Information Systems (ICIS 2019)*, 2019.
- [62] Matti Rossi, Christoph Mueller-Bloch, Jason Bennett Thatcher, and Roman Beck. Blockchain research in information systems: Current trends and an inclusive future research agenda. *Journal of the Association for Information Systems*, pages 1388–1403, 2019.
- [63] Jakob Roth, Fabian Schär, and Aljoscha Schöpfer. The tokenization of assets: Using blockchains for equity crowdfunding. *SSRN Electronic Journal*, 2019. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3443382.
- [64] Pablo Ruiz. Everything you always wanted to know about restricting token transfers but were afraid to ask. *Polymath Network*, 8.6.2018. URL <https://blog.polymath.network/all-you-ever-wanted-to-know-about-restricting-token-transfers-827009d649b7>.
- [65] Fabian Schär. Decentralized finance: On blockchain- and smart contract-based financial markets. *Review*, 103(2), 2021.
- [66] Joseph A. Schumpeter. *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Transaction Publishers, New Brunswick, Canada, 1934.

- [67] André Schweizer, Vincent Schlatt, Nils Urbach, and Gilbert Fridgen. Unchaining social businesses - blockchain as the basic technology of a crowdlending platform. *International Conference on Information Systems (ICIS 2017)*, 2017.
- [68] Armin Schwienbacher. Equity crowdfunding: anything to celebrate? *Venture Capital*, 21(1):65–74, 2019.
- [69] Joseph Stekli and Umit Cali. Potential impacts of blockchain based equity crowdfunding on the economic feasibility of offshore wind energy investments. *Journal of Renewable and Sustainable Energy*, 12(5):053307, 2020.
- [70] Abbey R. Stemler. The jobs act and crowdfunding: Harnessing the power—and money—of the masses. *Business Horizons*, 56(3):271–275, 2013.
- [71] Ali Sunyaev, Niclas Kannengießer, Roman Beck, Horst Treiblmaier, Mary Lacity, Johann Kranz, Gilbert Fridgen, Ulli Spankowski, and André Luckow. Token economy. *Business & Information Systems Engineering*, 2021.
- [72] Melanie Swan. *Blockchain: Blueprint for a new economy*. O’Reilly, Beijing, China, 2015.
- [73] Nick Szabo. Formalizing and securing relationships on public networks. *First Monday*, 2(9), 1997.
- [74] Jack Tanner. Summary of ethereum upgradeable smart contract r&d — part 1–2018. *Indorse*, 6.3.2018.
- [75] Nick Tomaino. On token value - the control. *The Control*, 6.8.2017.
- [76] Horst Treiblmaier, Melanie Swan, Primavera de Filippi, Mary Lacity, Thomas Hardjono, and Henry Kim. What’s next in blockchain research? *ACM SIGMIS Database: the DATABASE for Advances in Information Systems*, 52(1):27–52, 2021.
- [77] Karen E. Wilson and Marco Testoni. Improving the role of equity crowdfunding in europe’s capital markets. *SSRN Electronic Journal*, 2014. URL [10.2139/ssrn.2502280](https://ssrn.com/abstract=2502280).
- [78] Rebecca Yang, Ron Wakefield, Sainan Lyu, Sajani Jayasuriya, Fengling Han, Xun Yi, Xuechao Yang, Gayashan Amarasinghe, and Shiping Chen. Public and private blockchain in construction business process and information integration. *Automation in Construction*, 118:103276, 2020.

- [79] Dirk A. Zetsche, Ross P. Buckley, Douglas W. Arner, and Linus Föhr. The ico gold rush: It's a scam, it's a bubble, it's a super challenge for regulators. *SSRN Electronic Journal*, 2017. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3072298.
- [80] Zibin Zheng, Shaoan Xie, Hongning Dai, Xiangping Chen, and Huaimin Wang. Blockchain challenges and opportunities: a survey. *International Journal of Web and Grid Services*, 14(4), 2018.
- [81] Huasheng Zhu and Zach Zhizhong Zhou. Analysis and outlook of applications of blockchain technology to equity crowdfunding in china. *Financial Innovation*, 2(1), 2016.

Fundamentals of Perpetual Futures

Songrun He, Washington University of St. Louis
Asaf Manela, Washington University of St. Louis, Reichman
University
Omri Ross, University of Copenhagen, eToroX Labs
Victor von Wachter, University of Copenhagen

This paper is currently under peer review. This version is from April 2023.

Abstract Perpetual futures, swap contracts that never expire, are the most popular derivative traded in cryptocurrency markets, with more than \$100 billion traded daily. Perpetuals provide investors with leveraged exposure to cryptocurrencies, which does not require rollover or direct cryptocurrency holding. To keep the gap between perpetual futures and spot prices small, long position holders periodically pay short position holders a funding rate proportional to this gap. The funding rate incentivizes trades that tend to narrow the futures-spot gap. But unlike fixed-maturity futures, perpetuals are not guaranteed to converge to the spot price of their underlying asset at any time, and familiar no-arbitrage prices for perpetuals are not available, as the contracts have no expiry date to enforce arbitrage. Here, using a weaker notion of random-maturity arbitrage, we derive no-arbitrage prices for perpetual futures in frictionless markets and no-arbitrage bounds for markets with trading costs. These no-arbitrage prices provide a valuable benchmark for perpetual futures and simultaneously prescribe a strategy to exploit divergence from these fundamental values. Empirically, we find that deviations of crypto perpetual futures from no-arbitrage prices are considerably larger than those documented in traditional currency markets. These deviations comove across cryptocurrencies and diminish over time as crypto markets develop and become more efficient. A simple trading strategy generates large Sharpe ratios even for investors paying the highest trading costs on Binance, which is currently the largest crypto exchange by volume.

Keywords Perpetual futures, crypto, random-maturity, arbitrage, funding rate

Introduction

Perpetual futures are, by far, the most popular derivative traded in cryptocurrency markets, generating a daily volume of more than \$100 billion. Prior to the recent collapse of FTX, perpetual futures were among the most actively-traded products on the exchange, with the now-bankrupt hedge fund Alameda Research taking the other side of many such leveraged trades. Despite their central role in crypto markets, there is relatively little work studying these derivatives. In this paper we ask: what are the theoretical fundamental values of perpetual futures, and how large are deviations from these fundamentals empirically?

Perpetuals are a derivative that allows investors to speculate on or hedge against cryptocurrency price fluctuations using high leverage, without taking delivery of cryptocurrencies and without needing to roll them over. In essence, perpetuals are agreements between long and short counterparties. The party on the short side is required to pay the long side a sum based on the increase in the futures price between when they enter and exit the contract. Meanwhile, the long side pays the short side an ongoing cash flow called the ‘funding rate’. Different from traditional futures, perpetual futures do not have an expiration date. Both parties can enter or exit the contract at any time, and profits or losses are continually calculated and allocated to each side’s margin account. This setup improves the liquidity of the contract because there is no staggering of contracts with different maturities traded on the exchange and only a single perpetual futures contract per underlying is traded. Furthermore, this instrument does not require market participants to deal with rolling over of their futures positions and is traded 24/7.

Unlike fixed-maturity futures, perpetuals are not guaranteed to converge to the spot price of their underlying asset at any time, and familiar no-arbitrage prices for perpetuals are not available. To keep the gap between perpetual futures and spot prices small, long position holders periodically pay short position holders a funding rate proportional to this gap to incentivize trades that tend to narrow it. For example, when the futures price exceeds the spot price, arbitrageurs who borrow cash to long the spot and simultaneously short the futures would collect the funding rate. Their trades would tend to increase the spot price and decrease the futures price. A narrow gap means that perpetual futures provide effective exposure to variation in the spot price of the underlying asset to hedging and speculating investors. In practice, the funding rate is typically paid every eight hours and approximately equals the average futures-spot spread over the preceding eight hours. Note that the strategy just sketched, commonly referred to as ‘funding rate arbitrage’, is not

risk-free even if one ignores margin requirements and trading costs, simply because there is no predetermined maturity date when the trade would be unwound at a profit.

We derive no-arbitrage prices for perpetual futures in frictionless markets and derive no-arbitrage bounds in markets with trading costs. The theoretical perpetual futures price is proportional to the spot price of the underlying, with a constant of proportionality that increases in the ratio of the interest rate to the funding rate. The interest rate captures the cost of borrowing cash to finance holding the underlying, while the funding rate captures the benefit of shorting the futures. Thus, intuitively, the future-spot spread is larger when the cash borrowing interest rate is large relative to the funding rate.

Our derivation relies on a weaker notion of arbitrage that we call *random-maturity arbitrage*. As its name suggests, unlike traditional riskless arbitrage, we allow the strategy's time-to-maturity to be random. At first glance, one might object that such prices are not truly based on riskless arbitrage. Note, however, that riskless no-arbitrage pricing is usually just a useful fiction [23]. For example, in real-world futures markets, arbitrageurs must maintain a margin account during the entire period in which the arbitrage trade is open. Temporary worsening of apparent arbitrage opportunities can lead to liquidations and losses. As the saying goes, an arbitrageur must remain liquid longer than the market stays irrational. Thus, even arbitrage opportunities that appear to be riskless in theory, may be risky in practice.

These no-arbitrage prices provide a useful benchmark for perpetual futures and simultaneously prescribe a strategy to exploit divergence from these fundamental values. Motivated by the theoretical understanding, we study the empirical deviations of the perpetual futures price from the spot. The mean absolute futures-spot spread is about 60% to 100% per year across different cryptocurrencies, which is considerably larger than the deviations documented in traditional currency markets by [12]. We find strong comovement of the futures-spot gap across different cryptocurrencies. This comovement can be due to commonality in funding and market liquidity faced by arbitrageurs who operate in multiple cryptocurrencies. Common sentiment could also drive the difference in futures demand relative to the spot. Overall, the magnitude of the deviation is comparable to our theoretic no-arbitrage bound calibrated to actual trading fees.

The spread narrows considerably in 2022, suggesting a decrease in arbitrage frictions in the market and an increase in competition among arbitrageurs. The narrowing gap provides an additional perspective on the downfall of Alameda Research and Three Arrows Capital. According to news reports and interviews, both hedge funds seem to have pivoted from such arbitrage activity around late 2021 to early 2022 and started taking more directional bets on cryptocurrencies, with both direct unhedged crypto holdings and investments in crypto startups. The large declines in crypto prices in

2022 subsequently exhausted their capital and led to their bankruptcies^{1,2,3}. To understand the economics of the futures-spot spread, we consider a trading strategy motivated by the random-maturity arbitrage theory. Whenever the futures-spot spread exceeds the theoretical bound under certain trading cost tiers, we open the trading position and close it when the futures-spot spread returns to its theoretical relationship under no trading costs. We find that empirically, the random maturity arbitrage strategy generates a sizable Sharpe ratio even under high trading costs. For example, for Bitcoin perpetual futures, the strategy can generate a Sharpe ratio of 1.92 under high trading costs typical of retail investors, and up to 3.94 for highly-active market makers who pay no such fees. The performance is even better for ETH and other cryptocurrencies. The strategies deliver significant alphas relative to the 3-factor model of [20] and the 5-factor model of [9].

What explains these large no-arbitrage deviations? One natural explanation is that liquidity in crypto markets is insufficient for arbitrageurs to eliminate such violations. Our finding that the spreads decline over time is consistent with liquidity improving as these markets develop, and leaves open the possibility that they will narrow going forward. We also find, however, that past return momentum significantly explains the futures-spot gap with a time-series regression R^2 of more than 50%. When past returns are high, futures tend to be traded at a higher price relative to the spot. This indicates positive feedback or momentum trading behavior in the perpetual futures market. This correlation may linger even as crypto markets become more efficient.

The existing literature on perpetual futures mainly focuses on descriptive evidence. See e.g. [3], [24], [11], [14], and [26]. [5] provide a theoretical no-arbitrage analysis of the perpetual but they make over-simplifying assumptions by assuming the payoff from the perpetual is a fixed function of the underlying spot price. Compared to the existing literature, we illustrate the fundamental mechanism behind the perpetual design and derive theoretical no-arbitrage prices and bounds for this instrument.

Also related is recent literature on fixed maturity futures in the crypto space. [24] provides a comprehensive analysis of the carry of crypto futures, with the carry defined following the general definition of [18]. [24] document a volatile convenience yield in the crypto space driven by high leverage from trend-chasing small investors and the relative scarcity of arbitrage capital. [8]

¹See, for example, Forbes, November 19, 2022, on <https://www.forbes.com/sites/jeffkauffman/2022/11/19/how-did-sam-bankman-frieds-alameda-research-lose-so-much-money> How Did Sam Bankman-Fried's Alameda Research Lose So Much Money?

²Odd Lots, November 17, 2022, on <https://www.bloomberg.com/news/articles/2022-11-17/odd-lots-podcast-understanding-sam-bankman-fried-s-ftx-crypto-collapse> Understanding the Collapse of Sam Bankman-Fried's Crypto Empire

³Hugh Hendry's interview on December 3 2022 of Kyle Davies on the <https://www.youtube.com/watch?v=TzGdkB0xbCE> Collapse of Three Arrows Capital

provide a novel link of the volatile convenience yield to the staking, service flow, and transaction convenience of the underlying tokens. They show that the large deviation from uncovered interest rate parity can be reconciled with transaction convenience. Our paper focuses on perpetual futures rather than fixed maturity futures and extends fixed-maturity to random-maturity no-arbitrage pricing.

More broadly, our paper contributes to the understanding of the frictions and arbitrage in cryptocurrency markets. [22] study price deviations across exchanges. They find large gaps across countries, highlighting the important role played by capital controls and slow-moving arbitrage capital as in [13]. Our analysis focuses on the price wedge between the spot and the futures market. We find that even within an exchange, futures prices deviate from their theoretical arbitrage-free values. These results indicate there are significant limits to arbitrage as in [16] for cryptocurrencies in the early years.

The rest of the paper is organized as follows: Section 2 provides the institutional details and history of perpetual futures. Section 3 presents the no-arbitrage analysis of the perpetual futures market and derives the theoretical price of perpetual futures. Section 4 demonstrates the empirical futures-spot deviation and presents the simple theory-motivated trading strategy that can exploit the arbitrage opportunity. Section 5 provides some explanation for the deviation between futures and the spot. Section 6 concludes.

An Introduction to Perpetual Futures

The idea of perpetual futures was first introduced by [25]. The goal was to set up a perpetual claim on the cash flows of an illiquid asset. For example, the cash flow can be house rents and the underlying illiquid asset can be the real estate market. The purpose of the perpetual futures is to enable price discovery for the underlying with an illiquid or hard-to-measure price. Perpetual futures have no expiration date but cash is exchanged between the long and the short side: after buying the perpetual futures, the long side is entitled to receive the flow cash flow from the short side and they settle the price difference when exiting the position.

Perpetual futures in crypto markets similarly have no expiration date and cash is exchanged between the long and the short side, but their purpose is different from Shiller's original idea. First, unlike, e.g. real estate market, crypto has no inherent dividend or cash flow. Second, the price discovery argument of Shiller is most applicable to settings where spot prices are difficult to measure. Crypto spot prices, however, can be measured from active trading on different exchanges, and decentralized exchanges such as Bancor [17] or Uniswap [1] can offer price discovery for assets with minimal liquidity. The major role perpetual futures play in the crypto space is to offer an effective leveraged trading vehicle to hedge or speculate the underlying spot




Figure F.1: Total trading volumes of perpetual futures across exchanges

The figure displays the 7-day moving average daily traded volume for perpetual futures across all exchanges in blue. The median daily volume is \$17.8 bn. (2020), \$132.0 bn. (2021) and \$101.9 bn. (2022). This translates to a yearly volume of \$8,551 bn. (2020), \$51,989 bn. (2021) and \$39,306 bn. (2022) respectively. Additionally, the grey line depicts the ratio between the traded volume in perpetuals and spot markets. In 2022 the Perpetual markets are consistently trading between 2x or 3.5x the spot volume. The data is obtained from CoinGecko, a crypto data specialist. We exclude exchanges that are known for misrepresenting data (e.g. forms of wash trading).

price movement, which makes the market more complete. It also serves as an effective tax payment optimization tool for investors.

Crypto perpetual futures were first introduced by BitMEX in 2016, which gained great popularity in the crypto space since its inception. It initially served as an effective hedging tool for crypto miners. It was later adopted by crypto speculators interested in leveraged exposure. Nowadays, based on data from CoinGecko, the median total daily trading volume of perpetual futures across all exchanges is 101.9 billion in the year 2022 which is about $2\times$ to $3\times$ the total spot trading volume across these exchanges. Figure F.1 presents the 7-day moving average of the total trading volume of perpetual futures across all exchanges. We see a significant rise in trading of perpetual futures around January 2021 and the total volume stabilizes at a level above \$100 billion per day following the rise.

[10] document significant wash trading behavior among crypto exchanges because of competition, the ranking mechanism, and lack of regulation. They estimate that over 70% of crypto trading volume is not real. [4] confirm this conclusion using new data and an extended methodology. Therefore, in calculating trading volume, we exclude exchanges that are known to misrepresent data.

A key feature of crypto perpetual futures is the funding rate, which is the cash exchanged between the long and short counterparties. Its goal is to




Figure F.2: Bitcoin annualized funding rate across exchanges

The figure presents the 7-day moving average annualized funding rate for Bitcoin (BTC) across exchanges. The data is obtained from glassnode, an analytics platform. The data covers two major market turbulences: the COVID stock market crash from February to April 2020, and the FTX insolvency in November 2022. The FTX collapse led to significant negative funding rates on all solvent exchanges, thus the perpetual future prices were lower than the spot prices on the solvent exchanges. Vice versa on the insolvent FTX. The funding rates become more volatile after the event, indicating increased uncertainty for the market participants.

keep the futures price close to the underlying spot so that the futures can be an effective hedging tool for spot price movement. The funding rate is typically paid every 8 hours. Its value is approximately a weighted average of the prior 8 hours' price gap of the futures and the spot. If the futures price is above the spot, the funding rate will be positive, meaning the long side of the futures needs to pay the short side. This incentivizes traders to short the futures, and in doing so, to move its price back in line with the spot. On the other hand, when the futures price is below the spot, the funding rate turns negative, which means the short side needs to pay the long side. Therefore, the funding rate is the key mechanism for keeping the futures price close to the spot price.⁴ Note that perpetual futures cannot be replicated by rolling over 8-hour maturity futures. For the latter, every 8 hours, the price is guaranteed to converge to the underlying spot, while for the perpetual futures this is not the case.

Figure F.2 presents the annualized 7-day moving average of Bitcoin perpetual funding rates across leading exchanges from January 2020 to December 2022.

The funding rate is positive when the futures-to-spot spread is positive, and negative otherwise. Funding rates tend to be similar across exchanges due

⁴<https://www.binance.com/en/support/faq/360033525031> provides a detailed explanation of how the funding rate is calculated.

to cross-exchange arbitrage activity, but can diverge during extreme liquidity episodes. During March 2020, as Covid-19 started to spread and liquidity evaporated, funding rates turned substantially negative in most exchanges. Funding rates turned highly positive during the crypto bull run of early 2021. The last episode highlighted is the collapse of FTX, which was the 4th largest crypto exchange at the time. The figure shows that Bitcoin futures prices were substantially higher than spot prices at FTX, but the opposite was true on other exchanges. This pattern is consistent with FTX investors liquidating their short futures positions quickly, either voluntarily to reduce their exposure to the failing exchange, or involuntarily as the exchange liquidated their underfunded positions.

Arbitrage in Perpetual Futures

We next derive no-arbitrage prices for perpetual futures prices relative to spot prices. Unlike traditional futures, perpetual futures have no expiration date. To analyze arbitrage in this market, we extend the traditional notion of risk-free arbitrage where arbitrageurs have a guaranteed positive payoff at a certain time in the future. We first describe the payoff structure of perpetual futures. We then introduce a generalized notion of an arbitrage opportunity with a certain positive payoff but at an uncertain future time.

Definition 1. *A perpetual future $\{F_t\}_{t=0}^{\infty}$ written on $\{S_t\}_{t=0}^{\infty}$ is an agreement between the long and the short side. There is 0 cost to enter the agreement. After entering, both the long side and the short side can terminate the contract at any time t . Before termination, for each unit of perpetual future, the long must pay the short an \mathcal{F}_s -adapted cash-flow $\kappa(F_s - S_s)ds$, $s \in (0, t)$, referred to as the funding value. κ is a scaling parameter determining the magnitude of the funding rate relative to the price gap. At termination, the short needs to pay the long $F_t - F_0$ for each unit shorted.*

This definition is an approximation of real-world perpetual futures. In most exchanges, the funding value for perpetual futures is paid every 8 hours and approximately equals the difference between the futures price and the spot. If we measure time units in years, this setup would correspond to $\kappa = 1095$. Consider the case where the gap is constant $F_t - S_t = G$ over the 8-hour interval. The total funding rate payment over the 8 hours would be $\frac{\kappa G}{3 \times 365}$, which equals the price gap G in the empirical setting. Therefore, we have $\kappa = 1095$.

The traditional notions of arbitrage typically consider a guaranteed positive payoff at a certain future date.

Definition 2. *A (riskless) arbitrage opportunity is defined with respect to payoff x at a certain future time τ and its price $p(x)$. If the following conditions*

are satisfied: (1) $x \geq 0$ almost surely, (2) $x > 0$ with some positive probability, and (3) its price satisfies $p(x) \leq 0$, then this payoff is an arbitrage opportunity [7].

There is no expiration date in the perpetual futures market. Therefore, a generalized notion of arbitrage is required. Suppose the risk-free rate is constant. We define *random-maturity arbitrage* opportunities as zero-cost strategies with a guaranteed positive payoff at an uncertain future time. Stated formally:

Definition 3. *A random-maturity arbitrage opportunity is defined with respect to a random payoff x at a future random time $\tilde{\tau}$, $\tilde{\tau} \in (0, \infty)$, and its price $p(x)$. If the following conditions are satisfied: (1) $x \geq 0$ almost surely, (2) $x > 0$ with some positive probability, and (3) its price satisfies $p(x) \leq 0$, then this payoff is a random-maturity arbitrage opportunity.*

Definition 3 generalizes traditional arbitrage in the sense that there is a guaranteed positive payoff but at an uncertain future time. The following corollary specializes this definition for perpetual futures:

Corollary 1. *In the perpetual futures market, if a strategy (1) has 0 cost at time 0, and (2) for any price path of the futures and the spot, $\{F_t\}_{t=0}^{\infty}$ and $\{S_t\}_{t=0}^{\infty}$, there exists an unwinding time $\tilde{\tau}$ such that its discounted payoff at time $t = \tilde{\tau}$ is positive, then this strategy is a random-maturity arbitrage.*

We next show that when there is no random-maturity arbitrage, the gap between futures and spot prices is bounded by a constant. We make the two following assumptions:

Assumption 1. *The gap between the perpetual futures and the spot satisfies the following condition: $\liminf_{t \rightarrow \infty} |F_t(\omega) - S_t(\omega)| < \infty, \forall \omega$*

Assumption 2. *The risk-free rate r for arbitrageurs is constant.*

Assumption 1 is a no-bubble condition. \liminf represents the greatest lower bound as $t \rightarrow \infty$. In other words, we can always find F_t and S_t that lie within a finite bound as $t \rightarrow \infty$. This condition allows for the gap between F_t and S_t to explode in the limit as long as it shrinks at some subsequent time. Assumption 2 guarantees that there is no roll-over risk for the arbitrageurs. With these two assumptions, we can state the conclusion as follows.

Proposition 1 (No-arbitrage bound). *Under Assumptions 1 and 2, when there is a constant trading cost C when terminating the position, arbitrageurs will trade the perpetual futures until it lies within a bound of the spot:*

$$S_t \left(1 + \frac{r}{\kappa}\right) - C \leq F_t \leq S_t \left(1 + \frac{r}{\kappa}\right) + C . \quad (\text{F.1})$$

Proof. To prove this proposition, we consider the following two scenarios in Table F.1. (1) $F_0 > S_0(1 + \frac{r}{\kappa}) + C$; (2) $F_0 < S_0(1 + \frac{r}{\kappa}) - C$. We show in the first case, the arbitrageurs want to long the spot and short the futures because this is a random-maturity arbitrage opportunity and vice versa for the second case.

Table F.1 Discounted payoffs to arbitrage strategies in perpetual futures and spot markets

		$F_0 > S_0 \left(1 + \frac{r}{\kappa}\right) + C$	$F_0 < S_0 \left(1 + \frac{r}{\kappa}\right) - C$
Actions		Long spot, short futures	Long futures, short spot
Time 0	Futures	0	0
	Spot	$-S_0$	$+S_0$
	Cash	$+S_0$	$-S_0$
Time t	Futures	$F_0 - F_t$	$F_t - F_0$
	Spot	$+S_t$	$-S_t$
	Cash	$-S_0 e^{rt}$	$+S_0 e^{rt}$
	Funding	$\kappa \int_0^t (F_s - S_s) e^{r(t-s)} ds$	$-\kappa \int_0^t (F_s - S_s) e^{r(t-s)} ds$
	Trading Cost	$-C$	$-C$
Payoff		$e^{-rt} F_0 - e^{-rt} (F_t - S_t) - S_0 + \kappa \int_0^t (F_s - S_s) e^{-rs} ds - e^{-rt} C$	$e^{-rt} (F_t - S_t) - e^{-rt} F_0 + S_0 - \kappa \int_0^t (F_s - S_s) e^{-rs} ds - e^{-rt} C$

This table presents the costs and benefits of two arbitrage trading strategies: (1) when $F_0 > S_0 \left(1 + \frac{r}{\kappa}\right) + C$, long the spot and short the futures; (2) when $F_0 < S_0 \left(1 + \frac{r}{\kappa}\right) - C$, long the futures and short the spot. In the last row, the payoff from exiting the position is the discounted payoff from future and spot price changes, proceeds from the cash market, and the funding rate.

Scenario 1: If $F_0 > S_0 \left(1 + \frac{r}{\kappa}\right) + C$, consider the strategy of longing the spot and shorting the futures. We want to show that for any price path of the perpetual futures and the spot, there exists a future unwinding time t such that the strategy's payoff is positive, that is, this is a random-maturity arbitrage. Suppose to the contrary that there exists a price path $\{F_t\}_{t=0}^\infty$ and $\{S_t\}_{t=0}^\infty$ such that $\forall t$, the discounted payoff from the strategy is negative, or equivalently that:

$$\underbrace{e^{-rt} F_0 - S_0}_{\text{traditional spread}} + \underbrace{\kappa \int_0^t (F_s - S_s) e^{-rs} ds}_{\text{funding payments}} \leq \underbrace{e^{-rt} (F_t - S_t)}_{\text{spread at unwinding}} + \underbrace{e^{-rt} C}_{\text{trading cost}}. \quad (\text{F.2})$$

Denote $(F_t - S_t)e^{-rt} \equiv u_t$, the inequality changes into:

$$u_t \geq e^{-rt}(F_0 - C) - S_0 + \kappa \int_0^t u_s ds .$$

$\underline{u}_t = e^{-rt}(F_0 - C) - S_0 + \kappa \int_0^t \underline{u}_s ds$ provides a lower bound for all processes u_t satisfying the above inequality. Solving this integral equation, we have:

$$\underline{u}_t = \frac{F_0 r e^{-rt}}{\kappa + r} + \left(\frac{F_0 - C}{1 + \frac{r}{\kappa}} - S_0 \right) e^{\kappa t} .$$

When $F_0 > S_0 \left(1 + \frac{r}{\kappa}\right) + C$, $\lim_{t \rightarrow \infty} \underline{u}_t = \infty$. This contradicts Assumption 1, because it implies: $\liminf_{t \rightarrow \infty} (F_t - S_t)e^{-rt} \geq \lim_{t \rightarrow \infty} \underline{u}_t \rightarrow \infty$. Therefore, when $F_0 > S_0 \left(1 + \frac{r}{\kappa}\right) + C$, longing the spot and shorting the futures would be a random-maturity arbitrage opportunity.

Scenario 2: Next, if $F_0 < S_0 \left(1 + \frac{r}{\kappa}\right) - C$, consider the strategy of longing the futures and shorting the spot. Similarly, we want to show for this strategy, for any price path of the perpetual futures and the spot, there always exists a future time t such that the strategy payoff is positive, i.e. this is a random-maturity arbitrage. Suppose not, then there exists price path $\{F_t\}_{t=0}^\infty$ and $\{S_t\}_{t=0}^\infty$ such that $\forall t$, the discounted payoff is negative:

$$u_t \leq e^{-rt}(F_0 + C) - S_0 + \kappa \int_0^t u_s ds .$$

$\bar{u}_t = e^{-rt}(F_0 + C) - S_0 + \kappa \int_0^t \bar{u}_s ds$ provides an upper bound for all processes u_t satisfying the above inequality. This is the same integral equation as we see in the first case. We have:

$$\bar{u}_t = \frac{F_0 r e^{-rt}}{\kappa + r} + \left(\frac{F_0 + C}{1 + \frac{r}{\kappa}} - S_0 \right) e^{\kappa t} .$$

When $F_0 < S_0 \left(1 + \frac{r}{\kappa}\right) - C$, $\lim_{t \rightarrow \infty} \bar{u}_t = -\infty$. This contradicts Assumption 1, because $\liminf_{t \rightarrow \infty} (S_t - F_t)e^{-rt} \geq \lim_{t \rightarrow \infty} (-\bar{u}_t) \rightarrow \infty$. Therefore, when $F_0 < S_0 \left(1 + \frac{r}{\kappa}\right) - C$, longing futures and shorting the spot would be a random-maturity arbitrage opportunity. \square

No arbitrage prices are usually derived assuming away trading costs. For completeness, the following result considers this special case and provides the fundamental value of perpetual futures.

Proposition 2 (No-arbitrage price). *When there is no trading cost ($C = 0$), arbitrageurs will trade perpetual futures toward:*

$$F_t = S_t \left(1 + \frac{r}{\kappa}\right) . \quad (\text{F.3})$$

Proof. The proof follows that of Proposition 1 by setting $C = 0$. \square

To gain an intuition for this result consider the first three terms of Equation (F.2). The first is the traditional spread familiar from fixed-maturity futures pricing. Equating it to zero generates the usual no-arbitrage price for futures that pay no dividends and without carrying or storage costs. The second term is the present value of cumulated funding payments from initiation to unwinding. The first two terms are the gains to the arbitrageur in this scenario. The third term on the right-hand side of the inequality is the present value of the futures-to-spot spread at the random unwinding time and is the source of tension here. For fixed-maturity futures this gap is guaranteed to be zero at expiration, but an arbitrageur in perpetual futures faces the risk that when they wish to unwind the trade the spread explodes.

Our key insight is that a positive and even modestly large spread can still leave the arbitrageur with a positive payoff because funding payments also increase with the spread. As long as the spread does not diverge to infinity, there will be some future unwinding time t when the accumulated funding rate payments overcome any finite potential losses at unwinding. Arbitrage is absent when the accumulated benefits due to the funding rate κ exactly balance against the accumulated costs due to the borrowing interest rate r . The no-arbitrage price (F.3) captures this intuition and says that the futures-spot gap increases with the ratio r/κ .

Note that if we relax the assumption that the short rate is constant for arbitrageurs, this may introduce some additional risk to the payoff: (1) the arbitrageur will face some risk in rolling over her borrowing in the cash market; (2) she faces some risk in reinvesting the funding payment she receives. To map into the payoff in Table F.1, this corresponds to the cash market payoff of: $-S_0 e^{\int_0^t r_s ds}$ at time t and the funding payment of $\kappa \int_0^t (F_s - S_s) e^{\int_s^t r_u du} ds$, where $\{r_s\}_{s=0}^\infty$ is a stochastic process. Empirically, however, the volatility in the short rate is considerably smaller than the volatility in the gap between futures and spot prices. Moreover, the changes in the short rate have opposite effects on payment from the cash market and the funding payment. Therefore, we abstract from short-rate randomness here but note it is an interesting avenue for future work.

In the presence of trading costs, when the deviation of perpetual futures price from $S_0(1 + \frac{r}{\kappa})$, is larger than the round-trip trading costs C , arbitrageurs would have a strong incentive to trade perpetual futures toward the price $S_0(1 + \frac{r}{\kappa})$. This proposition also prescribes a trading strategy to exploit the futures-spot divergence in markets with different levels of trading costs. In the next part, we provide an empirical analysis of the futures-spot deviation and present arbitrage trading strategies motivated by our theory.

Data and Empirical Analysis

We conduct an empirical analysis of perpetual futures arbitrage strategies. We first describe the data. We then measure the deviations of the crypto futures-spot spread from the no-arbitrage benchmark. Finally, we implement a trading strategy that exploits deviations from random-maturity no-arbitrage bounds and quantify the gains from this strategy net of trading costs.

Data

We focus on the five largest cryptocurrencies excluding the stablecoins: Bitcoin (BTC), Ether (ETH), BNB (BNB), Dogecoin (DOGE), and Cardano (ADA) with a total market cap of \$529 billion, which account for 64.15% of the total market share of the Crypto market by November 2022.

For each token, we obtain perpetual futures and spot prices at a 1-hour frequency from Binance. Binance is by far the leading exchange in the crypto realm. Another major benefit of using the Binance data is: nearly every part of our trading strategies can be completed within the same platform without delay in transferring the fund. So it reflects the real-time investment opportunities facing the traders.

We also get the perpetual funding rate value from Binance. The funding rate is paid every 8 hours on Binance. So we have futures and spot prices every hour and realized funding rate payment at 8:00, 16:00, and 0:00 GMT each day. The perpetual and spot tradings are happening 24 hours per day and 7 days a week so there are no after-market hours in this market.

We get the earliest possible data on perpetual futures trading from Binance. The table below lists the starting and ending dates of our data for each crypto. Our data ends on 2022-11-13, covering the latest fallout of the FTX.

Table F.2 Sample descriptions

Crypto	Start date	End date	N
BTC	2019-09-10	2022-11-13	27,895
ETH	2019-11-27	2022-11-13	25,985
BNB	2020-02-10	2022-11-13	24,184
DOGE	2020-07-10	2022-11-13	20,559
ADA	2020-01-31	2022-11-13	24,424

This table presents the sample start and end dates for the 5 cryptocurrencies and their total number of observations.

We obtain the trading costs from Binance's website⁵. In general, the trad-

⁵<https://www.binance.com/en/fee/trading> provides data on trading fees in perpetual

ing costs for the spot market are significantly larger than that for the perpetual futures for similar trading volume because futures typically are traded with leverage. Fee tiers are attributed to the 30-day trading volume. We attribute high trading costs with a 30-day spot trading volume above \$1 million and futures trading volume above \$15 million (small individual trader). Medium trading costs attribute to a 30-day spot trading volume above \$150 million and futures trading volume above \$1 billion (small funds). Low trading costs attribute to a 30-day spot trading volume above \$2 billion and futures trading volume above 12.5 billion (large funds). No fee can be negotiated with customized contracts, for example for market makers. We consider trading costs for makers instead of takers because institutions typically trade maker orders.⁶ Table F.3 presents the specification of different trading costs. It also shows the random-maturity arbitrage bound for the deviation between the perpetual futures and the no-arbitrage price (ρ_l and ρ_u). The bounds become wider as the trading costs increase. Detailed explanations of the trading costs specifications are also provided in appendix Figure F.7 and Figure F.8.

To measure the deviation from the no-arbitrage price, we also obtain the interest rate data from Aave, a leading open-source DeFi liquidity protocol. Customers on Aave can either be a supplier or a borrower of cryptocurrencies. Because of the anonymity and decentralization of DeFi system, all borrowings are over-collateralized and the collateral can serve as an additional supply to the system for borrowing. The customers can also supply spare currencies to the system to earn an interest rate. The supply interest rate is typically different from the borrowing interest rate. Both interest rates and their wedge are determined algorithmically based on the market condition of supply and demand. We use interest rates from this platform because we believe it is a good proxy for the funding condition in the crypto market. Our results are robust if we use the interest rate from the traditional financial market.⁷

Our theory indicates using a risk-free rate available to the arbitrageurs. Therefore, we consider the interest rates on the stablecoins, which are not subject to volatile spot price movement and are the currency of denomination for perpetual futures margins. There are three major stablecoins traded on Aave: USDT, USDC, and DAI. To get a robust measure of the risk-free rate, we take an average of the three interest rates to arrive at our final risk-free supply and borrowing rate for the arbitrageurs.

We plot the time-series evolution of the interest rate in the appendix Figure F.10. During the early years, we see higher interest rate volatility due

futures and the spot market.

⁶Takers trade market orders while makers trade limit orders. Takers take liquidity from the market while makers make the market or provide liquidity to the market.

⁷We also run our analysis using daily T-bill rates obtained from Kenneth French's website at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_actors.html. The results are very close to ones using crypto market supply and borrowing rates.

to the funding liquidity of the DeFi platform. In later samples, interest rates are more stable and approach interest rates in traditional financial markets.

Our interest rate data starts from 2020-01-08. Therefore, for coins with perpetual data available before the time (BTC and ETH), we begin the analysis from 2020-01-08. For other coins (BNB, DOGE, ADA), we begin the analysis from the time when they have data available.

Table F.3 Trading costs specifications

Fee tier	Spot	Futures	ρ_l	ρ_u
No	0%	0%	0.0%	0.0%
Low	0.0225%	0.0018%	-53.2%	53.2%
Medium	0.045%	0.0072%	-114.4%	114.3%
High	0.0675%	0.0144%	-179.5%	179.2%

The different trading cost tiers: no, low, medium, and high. Fee tiers are assigned based on the past 30-day trading volume. High fees correspond to a 30-day trading volume above \$1mn in spot and above \$15mn in perpetuals, typically an individual trader. Medium fees attribute to a 30-day trading volume above \$150mn in spot and above \$1bn in perpetuals (small funds). Low fees attribute to a 30-day trading volume above \$2bn in spot and above \$12.5bn in perpetuals (large funds). The no fees tier can be negotiated with customized contracts, for example for market makers. We also report the theory implied no arbitrage bound (ρ_l and ρ_u) for ρ under different trading costs specifications. They are calculated using the following formulas: $\rho_l = \kappa \log(1 - C)$, $\rho_u = \kappa \log(1 + C)$, where C is the round-trip percentage trading costs of the long-short strategy of perpetual and spot.

Deviations of Perpetual Futures from No-arbitrage Benchmarks

The focus of our empirical analysis is the annualized deviation ρ , defined as the interest rate that rationalizes an observed future-spot spread:

$$F = S \left(1 + \frac{r + \rho}{\kappa} \right) .$$

Using f and s to denote $\log(F)$ and $\log(S)$ respectively, we obtain the following approximate equation for ρ :

$$\rho \approx \kappa(f - s) - r . \tag{F.4}$$

This definition is the same in spirit to [12] who define the CIP deviation as the wedge that would equate the dollar borrowing rate and the synthetic dollar borrowing rate.

At each hour, we calculate ρ using data on the perpetual futures price, spot price, and the crypto risk-free interest rate. When the perpetual futures price is above the spot, an arbitrageur would short the futures and long the spot, and would finance her position by borrowing in the cash market (as is shown in Table F.1). Therefore, we use the borrowing rate from Aave as the risk-free rate. On the other hand, when the perpetual futures price is below the spot, an arbitrageur would long the futures, short the spot, and invest the proceeds from shorting in the cash market (as is shown in Table F.1). In such cases, we use the supply rate from Aave as the risk-free rate.




Figure F.3: Deviations of perpetual futures from no-arbitrage benchmarks

This figure presents the 7-day moving averages of the annualized deviation of the perpetual futures-spot spread from the no-arbitrage benchmark for the five cryptocurrencies. The deviation ρ is defined as $\rho = \kappa(f - s) - r$.

Figure F.3 presents the 7-day moving average of ρ for each of the five cryptocurrencies. There is significant comovement in futures-spot spreads across all five cryptocurrencies, suggesting there exists a common factor driving the wedge. Figure F.4 confirms this finding: the futures-spot deviations are highly correlated across cryptocurrencies. This phenomenon highlights the integration of the cryptocurrency markets. One hypothesis is that the comovement in ρ across various cryptocurrencies reflects time-varying funding constraints experienced by arbitrageurs in the market, as illustrated in [6] and [15]. Since arbitrageurs are marginal traders in all markets, their funding constraints are manifested in the ρ across all cryptocurrency markets. Our theoretical and empirical analyses emphasize the importance of examining the object of ρ , as it sheds light on market stress and offers an estimate of the shadow costs associated with arbitrageurs' trading and funding constraints.

On the other hand, from the demand side, the commonality in sentiment




Figure F.4: Correlation of ρ , funding rate, spot returns and fear & greed index

This figure presents the correlation among different crypto’s futures-spot deviation ρ (rows 1 to 5), funding rate (rows 6 to 10), fear & greed index (row 11) and the spot returns (rows 12 to 16).

across different cryptocurrencies can also contribute to explaining the comovement of ρ across different markets. From our theoretical analysis, arbitrageurs will only accommodate the demand in the market if the price deviation exceeds the trading and funding costs. Consequently, the overall sentiment in the futures market relative to the spot market manifests in the price gap between futures and spot. If the sentiment in different markets is driven by a common factor, we expect to observe a high comovement in the futures-spot spread. In Section 5, we provide further analysis to explain the futures-spot spread. Our findings reveal that past returns of each cryptocurrency serve as significant explanatory variables for the time-series variation in the spread.

Interestingly, the correlation between deviations ρ and spot market returns is low. Although spot market returns are highly correlated among themselves, in line with the strong market factor results from [20], and no-arbitrage price deviations are highly correlated among themselves, it appears that different forces drive these distinct phenomena.

Furthermore, we find that after the year 2022, futures-spot spreads become smaller in magnitude and less volatile compared to earlier years. The 7-day moving average stays around -50% most of the time for the 5 cryptocurrencies while larger sways in earlier years are quite common. This suggests the market is becoming increasingly efficient. In terms of the level of the deviation, which

appears to stabilize in the negative region, there can be two forces: (1) on the futures customer end, the relative end demand in the futures market is weaker compared to the spot; (2) for arbitrageurs, because of the lack of infrastructure to short the cryptos in the spot market (high shorting costs), their funding constraints would be larger in the negative region. All these forces contribute to the stabilization of the futures spot deviation around -50 percent.

Last but not least, by design, ρ is also highly correlated with the funding rate. The funding rate does not correlate perfectly with ρ because in real-world implementations: (1) there is a clamp region, within which the funding rate equals 0.01% ; (2) the funding rate is calculated as a weighted average of futures-spot price deviations, with larger weight given to more recent observations; (3) in calculating the funding rate, Binance does not just consider the quote price, they also use an impact margin notional to consider price impact of the trade.⁸

Random-maturity Arbitrage Strategy

In this part, we provide a trading strategy motivated by our random-maturity arbitrage theory. Table F.3 reports for different trading cost tiers, the bounds (ρ_l and ρ_u) beyond which there exist random-maturity arbitrage opportunities. We consider a simple trading strategy: whenever ρ enters the region outside the annualized round-trip trading costs in Table F.3, we open the position. We close the position when ρ first goes back to 0.

Figure F.5 provides an illustration of the trading strategy. When the deviation is beyond the orange lines, the strategy opens a trading position. The position is closed when the futures/spot deviation first hits the red line. We present trading thresholds across different trading costs for different currencies in Figure F.11 in the appendix. Different trading strategies have the same close-position line which is equal to the risk-free rate while the open-position line adjusts to the level of trading costs.

Since our trading strategy is a threshold trading rule, to annualize Sharpe ratios we follow [21], which considers a trading strategy that is only active on FOMC announcement dates. We first calculate the mean (μ) and standard deviation (σ) of our trading strategy during the time it is active. Next, we scale μ/σ by the number of periods the strategy is active in a year:

$$SR = \frac{\mu}{\sigma} \sqrt{N_a} ,$$

where μ and σ are average hourly returns and N_a is the average number of hours the strategy is active in a year. We follow the same approach to annualize the returns and standard deviations: $\mu_{ann} = \mu N_a$, $\sigma_{ann} = \sigma \sqrt{N_a}$.

⁸See <https://www.binance.com/en/support/faq/360033525031> for details of the funding rate calculation.




Figure F.5: Random-maturity Arbitrage Strategy: Bitcoin, high trading costs
 Futures-spot deviations and trading thresholds of the random-maturity arbitrage strategy we implement for BTC under high trading costs. Each blue dot in the figure represents the annualized deviation of futures from the spot $\rho = \kappa(f - s) - r$. The orange line is the open-position threshold and the red line is the close-position threshold.

Table F.4 presents the strategy's performance under different fee tiers. As trading costs decrease, the random-maturity arbitrage strategy engages in more active trading, leading to an increase in the SRs and a decrease in the average duration of open-to-close positions. The strategy also delivers highly significant alphas relative to the 3-factor model by [20] and the 5-factor model by [8]. Evidently, the strategy's performance cannot be explained by previously suggested risk factors.

Table F.5 zooms in and presents the trading performance of the strategy under high trading costs. We consider two cases: (1) unrestricted; (2) long-spot only. We consider the second case because, in earlier years of crypto derivatives, the infrastructure of shorting cryptocurrencies is not well-developed. We find that, as is implied by our theory, whenever the deviation between futures and spot is larger than the trading costs, performing a random-maturity arbitrage strategy would generate returns with high SR. After trading costs, the strategy generates a SR of 1.92 for BTC and much higher SRs for other cryptocurrencies.

Table F.5 also shows that as time goes by, the perpetual futures market seems to become more and more efficient. As we can see in 2022, the deviation of crypto price from the arbitrage-free bound is much less frequent compared to earlier years. But when the deviation happens, the resulting SR from the trade remains high.

Table F.4 Performance of Random-maturity Arbitrage Strategy

		Fee tiers			
		No	Low	Medium	High
BTC	SR	3.94	2.50	2.31	1.92
	Return	17.89	11.21	10.13	8.15
	Volatility	4.54	4.48	4.39	4.25
	MaxDD	-4.24	-4.27	-4.34	-4.43
	α	22.64	10.86	8.24	5.47
	t_α	5.39	2.71	2.10	1.38
	Active %	100.00	84.93	43.87	22.32
	OtC time	15.80	63.90	95.97	113.38
ETH	SR	5.43	3.46	3.13	2.82
	Return	28.03	17.49	15.17	12.68
	Volatility	5.16	5.06	4.85	4.49
	MaxDD	-4.13	-4.21	-3.90	-3.94
	α	45.55	23.09	18.33	14.42
	t_α	7.56	4.25	3.69	3.15
	Active %	99.94	84.63	51.21	28.11
	OtC time	12.10	37.57	70.75	75.77
BNB	SR	12.49	8.07	6.41	5.44
	Return	62.38	38.40	29.14	23.43
	Volatility	4.99	4.76	4.55	4.31
	MaxDD	-1.12	-1.12	-1.10	-1.11
	α	107.81	58.17	41.12	32.87
	t_α	10.53	7.79	7.08	6.50
	Active %	99.97	89.25	66.67	39.94
	OtC time	7.87	14.65	22.20	24.07
DOGE	SR	12.61	8.34	5.96	4.42
	Return	112.09	70.88	49.74	35.82
	Volatility	8.89	8.49	8.35	8.10
	MaxDD	-8.56	-8.56	-8.56	-8.56
	α	312.42	151.20	88.73	54.65
	t_α	7.76	6.56	5.53	4.46
	Active %	99.85	90.35	68.52	36.90
	OtC time	5.56	8.50	11.21	11.26
ADA	SR	11.34	6.46	4.07	3.12
	Return	69.58	38.05	23.33	16.72
	Volatility	6.13	5.89	5.73	5.36
	MaxDD	-4.30	-4.26	-4.29	-4.34
	α	120.52	52.28	27.52	18.11
	t_α	14.56	10.22	6.55	4.59
	Active %	99.64	90.49	69.32	41.03
	OtC time	6.57	11.26	20.80	33.77

Performance under different trading cost tiers. The fees for spot (futures) are 2.25 (0.18) bps, 4.5 (0.72) bps, and 6.75 (1.44) bps for the low, medium, and high trading cost levels. Statistics reported are the annualized Sharpe ratio, return (%), volatility (%), max drawdown (%), alpha (%), t-stat of the alpha, the proportion of time the strategy is active (%), and average open-to-close (OtC) position duration in hours.

Table F.5 Performance Over Time: High Trading Costs Tier

		Unrestricted				Long-spot Only			
		2020	2021	2022	All	2020	2021	2022	All
BTC	SR	2.26	2.39	1.23	1.92	1.98	2.28	0.55	1.75
	Return	8.28	14.80	0.26	8.15	6.43	13.96	0.10	7.16
	Volatility	3.67	6.18	0.21	4.25	3.25	6.12	0.18	4.10
	Active %	28.66	34.43	1.01	22.32	22.18	32.07	0.03	18.95
	N	8,616	8,760	7,536	24,912	8,616	8,760	7,536	24,912
ETH	SR	3.08	3.52	1.39	2.82	2.57	3.37	0.87	2.50
	Return	17.11	18.08	1.36	12.68	13.99	17.14	0.62	11.06
	Volatility	5.55	5.13	0.98	4.49	5.44	5.09	0.71	4.42
	Active %	37.80	34.91	9.12	28.11	34.40	34.54	0.09	24.07
	N	8,616	8,760	7,536	24,912	8,616	8,760	7,536	24,912
BNB	SR	6.04	6.60	2.61	5.44	4.45	5.00	0.68	3.94
	Return	31.26	33.13	4.04	23.43	15.32	21.94	0.16	12.99
	Volatility	5.17	5.02	1.55	4.31	3.44	4.39	0.24	3.29
	Active %	55.07	48.03	14.84	39.94	34.96	31.27	0.03	22.70
	N	7,815	8,760	7,536	24,111	7,815	8,760	7,536	24,111
DOGE	SR	4.90	5.93	1.93	4.42	2.64	5.05	1.11	3.10
	Return	59.81	53.51	1.94	35.82	31.16	36.34	0.29	22.02
	Volatility	12.19	9.02	1.00	8.10	11.80	7.19	0.26	7.11
	Active %	60.12	49.19	9.70	36.90	37.49	41.91	0.05	25.61
	N	4,190	8,760	7,536	20,486	4,190	8,760	7,536	20,486
ADA	SR	3.87	3.41	2.27	3.12	3.16	2.94	0.00	2.49
	Return	21.32	24.18	3.15	16.72	16.03	20.36	0.00	12.62
	Volatility	5.51	7.09	1.39	5.36	5.08	6.92	0.00	5.08
	Active %	47.61	42.18	32.66	41.03	36.41	39.53	0.00	26.27
	N	8,055	8,760	7,536	24,351	8,055	8,760	7,536	24,351

This table presents the Sharpe ratios, annualized returns (%), standard deviations (%), and active percentages (%) of the random-maturity arbitrage trading strategies for five different cryptocurrencies with high trading costs. We break down returns for each year and provide summary statistics across all time. The left panel shows the performance of the unrestricted trading strategy, where both long and short spot positions are allowed. The right panel shows the performance of the long-spot-only trading strategy, where shorting the spot is not allowed.

Comparing the results with [12], we find that deviations in crypto perpetual futures are considerably larger in magnitude. As a result, gains from the arbitrage strategies we study are also larger. Even though the volatility of the trading strategy also scales up, the Sharpe ratios in the crypto space are still larger than those in the traditional foreign exchange market as reported in [12].

When a futures-spot gap opens up, gains from the trading strategy could potentially arise from two main sources: price convergence and funding rate payments. While industry publications usually emphasize the funding rate channel, we note that price convergence can generate quicker gains from arbi-

trage if dislocations are short-lived. To examine these two sources empirically, in Table F.6 we provide a decomposition of the trading strategy's performance into price convergence versus funding rate payment. We find that price convergence plays a dominant role in total trading returns, while funding rate payments have a more minor role, which seems to diminish over time.

Table F.6 Return Decomposition: Price Convergence vs Funding Rate Payment for Random-maturity Arbitrage Strategies

		2020	2021	2022	All
BTC	Return	21.68	29.60	-0.05	17.89
	Price	14.00	14.29	2.11	10.51
	Funding	7.68	15.31	-2.16	7.39
ETH	Return	40.93	36.66	3.24	28.03
	Price	27.58	19.07	3.75	17.38
	Funding	13.35	17.59	-0.51	10.65
BNB	Return	82.48	76.29	25.36	62.38
	Price	66.70	60.37	20.00	49.80
	Funding	15.78	15.92	5.36	12.57
DOGE	Return	220.81	113.79	49.67	112.09
	Price	214.72	93.24	50.52	102.37
	Funding	6.09	20.55	-0.85	9.72
ADA	Return	90.81	66.73	50.19	69.58
	Price	77.17	49.51	49.05	58.52
	Funding	13.64	17.22	1.13	11.06

This table decomposes the portfolio return into the part due to price convergence and the part due to funding rate payment.

The success of the trading strategy supports the theory of random-maturity arbitrage. Whenever there is a deviation larger than the gap implied by the theory, betting on convergence tends to generate a positive payoff at some uncertain future time. Therefore, the convergence arbitrage trade generates high Sharpe ratios.

Explaining Futures-spot Deviations

From Figure F.3, we observe a strong common comovement of the futures-spot deviation across all crypto assets. Our goal is to understand the fundamental forces driving this common factor. We consider two potential hypotheses: (1) the time-varying funding constraints of arbitrageurs, and (2) the time-varying relative demand from end-users for perpetual futures compared to the spot.

Both forces could potentially explain these patterns. Arbitrageurs will accommodate the relative demand from end-users only until the price deviation

lies within the random-maturity no-arbitrage bound, which depends on arbitrageurs' funding conditions. As arbitrageurs are likely the marginal investors in all cryptocurrency markets, their time-varying funding constraints may create common time-series variation in ρ across different cryptocurrencies.

On the other hand, when the deviation lies within the no-arbitrage bound, variations in demand from the perpetual futures market compared to the spot market will influence the price deviation. It is plausible that relative demand has common factors across different cryptocurrencies. For instance, sentiment could drive their common variation, as perpetual futures allow for high leverage, which attracts overconfident, extrapolative, and sentiment-driven investors.

Determining which of the two factors better explains the observed price deviation is an empirical question. To shed light on this, we use past returns as a proxy for relative extrapolative demand in the perpetual futures market compared to the spot market. Past returns correlate with the demand from investors with extrapolative beliefs, who are more likely to trade in the perpetual futures market given the high leverage it provides. Therefore, past returns would correlate with the relative demand in the perpetual futures market compared to the spot.

For the time-varying funding constraint, we consider using crypto return volatility as a proxy, because, as in [2], arbitrageurs likely face Value-at-Risk (VaR) type constraints, which are more likely to bind when market volatility is high.

Table F.7 presents the regression results of ρ on past returns, volatility, and both for BTC and ETH. The regression coefficients on past returns are positive and highly significant. This suggests that when past returns are high, perpetual futures exhibit a more positive deviation against the spot. Mapping this back to our second hypothesis, when past returns are high, the demand for futures relative to the spot is also likely to be high. Consequently, even after arbitrageurs accommodate the demand outside of trading costs, the residual demand still manifests itself through the perpetual-spot deviation.

This general observation on limits to arbitrage is also connected to the findings of [22], who discovered that crypto price deviations across international exchanges tend to comove with one another. The driving force behind this comovement is investors' buying pressure across various countries, with cross-country capital control serving as the primary impediment to arbitrage capital. In the context of perpetual futures, end-user demand contributes to the comovement of different cryptocurrencies. The limits to arbitrage predominantly manifest in the form of trading frictions.

[19] document a significant time-series momentum pattern in the crypto market. Given that positive past returns lead to a positive gap between futures and spot prices, it is worth examining further whether the time-series momentum phenomenon is driven by margin trading and price pressure from the perpetual futures market. We leave this as an open question for future

Table F.7 Regression of the futures-spot gap against explanatory variables

	BTC			ETH		
Ret	0.28*** (7.79)		0.28*** (7.74)	0.20*** (3.87)		0.23*** (4.83)
Vol		-0.02 (-1.60)	-0.01 (-0.64)		-0.01** (-1.99)	0.01 (1.58)
Const	-0.12** (-2.16)	0.30 (1.22)	-0.03 (-0.16)	-0.06 (-0.58)	0.36 (1.40)	-0.28* (-1.67)
R2	0.55	0.03	0.56	0.47	0.03	0.49
N	1011	1011	1011	1011	1011	1011

We present regression results of ρ on the past four months' annualized returns (Ret) and volatility (Vol) for Bitcoin and Ethereum. The left panel reports results for Bitcoin, and the right panel reports results for Ethereum. For each cryptocurrency, we consider three models: (1) regressing ρ on past returns; (2) regressing ρ on past volatility; and (3) regressing ρ on both past returns and volatility. The observations are at a daily frequency. The sample period spans from January 30, 2020, to November 13, 2022, totaling 1,011 days. The table also reports the R-squared values for each regression model. We report HAC-robust t-statistics in parentheses. $*p < .1$, $**p < .05$, $***p < .01$

research.

We find that volatility does not significantly covary with the futures and spot deviation. This suggests that it is not the time-varying funding constraints of arbitrageurs that drive the comovement in futures-spot deviations across different cryptocurrencies.

In summary, a prerequisite for the gap between futures and spot prices to occur is the existence of trading costs for arbitrageurs. Arbitrage trading will accommodate all the demand in the futures until the price deviation lies within a trading cost bound. Within the bound, the relative demand of futures compared to the spot will still manifest itself in the deviation of futures prices from the spot. Due to the comovement of the time-varying relative demand, we observe a significant common factor in crypto future-spot deviations.

Conclusion

Perpetual futures play an important role in today's crypto markets and could potentially be adopted in non-crypto markets in the future. Understanding the fundamental mechanism of this financial derivative is a crucial first step for understanding speculation and hedging dynamics in this fast-evolving area. We provide a comprehensive analysis of the arbitrage and funding rate payment mechanisms that underpin perpetual futures.

In an ideal, frictionless world, we show that arbitrageurs would trade perpetual futures in such a way that a constant relationship exists between the future price and the spot price. In the presence of trading costs, the deviation of the futures price to the spot would lie within a bound.

Motivated by the theory, we then empirically examine the comovement of the futures-spot spread across different cryptocurrencies and implement a theory-motivated arbitrage strategy. We find that this simple strategy yields substantial Sharpe ratios across various trading cost scenarios. The evidence supports our theoretical argument that perpetual futures-spot spreads exceeding trading costs represent a random-maturity arbitrage opportunity.

Finally, we provide an explanation for the common comovement in futures-spot spreads across different crypto-currencies: arbitrageurs can only accommodate market demand if the price deviation exceeds trading costs. As a result, the overall sentiment in the futures market relative to the spot market is reflected in the spread. Our empirical findings suggest that past return momentum can account for a significant portion of the time-series variation in the futures-spot spread.

Appendix

A. Additional Figures and Tables

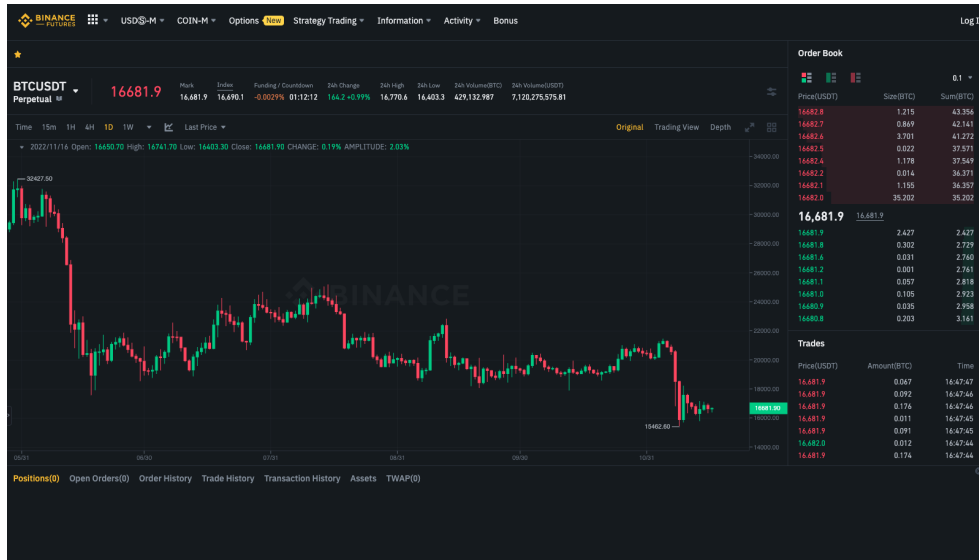


Figure F.6: Trading view of BTC perpetual futures on Binance

This figure presents the perpetual futures trading view on Binance. The key information includes futures price (Mark), spot price (Index), real-time funding rate based on the rolling average of the past 8 hours' observations, and countdown toward funding rate payment (Funding / Countdown) (illustrated with a red rectangle). In this example, futures are trading at a lower price than the spot. The funding rate is negative and is to be paid in 1 hour and 12 minutes.

Fee Rate

Spot Trading Margin Borrow Interest USD \mathbb{S} -M Futures Trading COIN-M Futures Trading Options Trading Swap Farming P2P

Level	30d Trade Volume (BUSD)	and/or	BNB Balance	Maker / Taker	Maker / Taker BNB 25% off
Regular User	< 1,000,000 BUSD	or	≥ 0 BNB	0.1000% / 0.1000%	0.0750% / 0.0750%
VIP 1	≥ 1,000,000 BUSD	and	≥ 25 BNB	0.0900% / 0.1000%	0.0675% / 0.0750%
VIP 2	≥ 5,000,000 BUSD	and	≥ 100 BNB	0.0800% / 0.1000%	0.0600% / 0.0750%
VIP 3	≥ 20,000,000 BUSD	and	≥ 250 BNB	0.0700% / 0.1000%	0.0525% / 0.0750%
VIP 4 🌟	≥ 100,000,000 BUSD	and	≥ 500 BNB	0.0200% / 0.0400% 0.0700% / 0.0900%	0.0150% / 0.0300% 0.0525% / 0.0675%
VIP 5 🌟	≥ 150,000,000 BUSD	and	≥ 1,000 BNB	0.0200% / 0.0400% 0.0600% / 0.0800%	0.0150% / 0.0300% 0.0450% / 0.0600%
VIP 6 🌟	≥ 400,000,000 BUSD	and	≥ 1,750 BNB	0.0200% / 0.0400% 0.0500% / 0.0700%	0.0150% / 0.0300% 0.0375% / 0.0525%
VIP 7 🌟	≥ 800,000,000 BUSD	and	≥ 3,000 BNB	0.0200% / 0.0400% 0.0400% / 0.0600%	0.0150% / 0.0300% 0.0300% / 0.0450%
VIP 8 🌟	≥ 2,000,000,000 BUSD	and	≥ 4,500 BNB	0.0200% / 0.0400% 0.0300% / 0.0500%	0.0150% / 0.0300% 0.0225% / 0.0375%
VIP 9	≥ 4,000,000,000 BUSD	and	≥ 5,500 BNB	0.0200% / 0.0400%	0.0150% / 0.0300%

- "Taker" is an order that trades at a market price, "Maker" is an order that trades at a limited price. [Learn more](#)
- VIP trade volume levels are measured on the basis of the spot trading volume, or whether the futures trading volume meets the standard (Futures trading volume includes USD \mathbb{S} -M futures and COIN-M futures).
- Refer friends to earn trading fees 20% kickback. [Learn more](#)

Figure F.7: Trading costs tiers for the spot market

This figure presents the trading costs tiers for the crypto spot market from Binance: <https://www.binance.com/en/fee/trading>. Our high, medium, and low trading cost specification corresponds to VIP 1, VIP 5, and VIP 8 tiers as illustrated with red rectangles in the picture. The 30-day trading volume requirements for VIP 1, 5, and 8 are 1 million, 150 million, and 2 billion respectively in the spot market. We consider the trading costs for makers as institutions typically trade maker orders. Binance offers a temporary discount for VIP 4-8 to have the same trading cost as VIP 9. We consider the non-discounted trading costs to make the comparison more reliable and fair.

B. Data-driven Arbitrage Strategy

In the main text of our paper, we demonstrate the profitability of a simple theory-motivated trading strategy. The trading threshold can also be potentially further improved using a data-driven approach. In this part, we implement a data-driven two-threshold arbitrage trading strategy in the per-

Fee Rate

Spot Trading Margin Borrow Interest **USDS-M Futures Trading** COIN-M Futures Trading Options Trading Swap Farming P2P

Level	30d Trade Volume (BUSD)	and/or	BNB Balance	USDT Maker / Taker	USDT Maker/Taker BNB 10% off	BUSD Maker / Taker	BUSD Maker/Taker BNB 10% off
Regular User	< 15,000,000 BUSD	or	≥ 0 BNB	0.0200%/0.0400%	0.0180%/0.0360%	0.0120%/0.0300%	0.0108%/0.0270%
VIP 1	≥ 15,000,000 BUSD	and	≥ 25 BNB	0.0160%/0.0400%	0.0144%/0.0360%	0.0120%/0.0300%	0.0108%/0.0270%
VIP 2	≥ 50,000,000 BUSD	and	≥ 100 BNB	0.0140%/0.0350%	0.0126%/0.0315%	0.0120%/0.0300%	0.0108%/0.0270%
VIP 3	≥ 100,000,000 BUSD	and	≥ 250 BNB	0.0120%/0.0320%	0.0108%/0.0288%	0.0120%/0.0300%	0.0108%/0.0270%
VIP 4	≥ 600,000,000 BUSD	and	≥ 500 BNB	0.0100%/0.0300%	0.0090%/0.0270%	0.0100%/0.0300%	0.0090%/0.0270%
VIP 5	≥ 1,000,000,000 BUSD	and	≥ 1,000 BNB	0.0080%/0.0270%	0.0072%/0.0243%	-0.0100%/0.0230%	-0.0100%/0.0207%
VIP 6	≥ 2,500,000,000 BUSD	and	≥ 1,750 BNB	0.0060%/0.0250%	0.0054%/0.0225%	-0.0100%/0.0230%	-0.0100%/0.0207%
VIP 7	≥ 5,000,000,000 BUSD	and	≥ 3,000 BNB	0.0040%/0.0220%	0.0036%/0.0198%	-0.0100%/0.0230%	-0.0100%/0.0207%
VIP 8	≥ 12,500,000,000 BUSD	and	≥ 4,500 BNB	0.0020%/0.0200%	0.0018%/0.0180%	-0.0100%/0.0230%	-0.0100%/0.0207%
VIP 9	≥ 25,000,000,000 BUSD	and	≥ 5,500 BNB	0.0000%/0.0170%	0.0000%/0.0153%	-0.0100%/0.0230%	-0.0100%/0.0207%

• "Taker" is an order that trades at a market price, "Maker" is an order that trades at a limited price. [Learn more](#)

• VIP trade volume levels are measured on the basis of the spot trading volume, or whether the futures trading volume meets the standard (Futures trading volume includes USDS-M futures and COIN-M futures).

• Refer friends to earn trading fees 20% kickback. [Learn more](#)

Figure F.8: Trading costs tiers for the perp market

This figure presents the trading costs tiers for the crypto perpetual market from Binance: <https://www.binance.com/en/fee/trading>. Our high, medium, and low trading cost specification corresponds to VIP 1, VIP 5, and VIP 8 tiers as illustrated with red rectangles in the picture. The 30-day trading volume requirements for VIP 1, 5, and 8 are 15 million, 1 billion, and 12.5 billion respectively in the perpetual market. We consider the trading costs for makers as institutions typically trade maker orders.

petual futures market. The strategy can be characterized with a tuple of two thresholds: (u, l) , where u denotes the upper bar and l denotes the lower bar, $u > l$. When $\rho > u$, we long the spot and short the futures. When $-l < \rho < l$, we close the position. When $\rho < -u$, we long the futures and short the spot. Figure F.16 presents an illustration of the strategy for Bitcoin with the trading thresholds estimated using real-time data.

To determine the optimal (u, l) , at the beginning of each month, we calculate the returns of the two-threshold trading strategy based on the past 6 months' data on a grid of parameters. From Table F.3, we find the theory implied bounds for price deviation across all trading costs specifications lie within -200% to $+200\%$. So we choose the grid as increasing from 0% to




Figure F.9: Ether annualized funding rate across exchanges

The figure presents the 7-day moving average annualized funding rate for Ether (ETH) across exchanges. The data is obtained from glassnode, an analytics platform. The data covers two major market turbulences: the COVID stock market crash from February to April 2020, and the FTX insolvency in November 2022. The FTX collapse led to significant negative funding rates on all solvent exchanges, thus the perpetual future prices were lower than the spot prices on the solvent exchanges. Vice versa on the insolvent FTX. The funding rates become more volatile after the event, indicating increased uncertainty for the market participants.

200% with an incremental step of 10% for u and l ($u \geq l$). In total, there are 210 model specifications. We choose the model that delivers the highest Sharpe Ratio in the validation sample of the past 6 months.

Figure F.16 presents the real-time trading thresholds of our arbitrage strategy for BTC under high trading costs. To mitigate the trading cost, the strategy automatically chooses a much lower close-position threshold compared to the open-position threshold. We also present the visualization of all trading strategies under different trading costs in appendix Figure F.12 to F.15. Under no-trading costs, the strategy chooses a low open-position threshold and the open- and close-position thresholds coincide with each other. Since there is no trading cost, the strategy no longer needs to wait for price convergence to avoid high turnover. On the contrary, with high trading costs, the strategy chooses a larger open-position threshold and a smaller close-position threshold. This reflects the automatic adjustment from the algorithm for trading costs and turnovers.

Table F.8 presents the return statistics for this strategy under high trading costs over time. In the baseline ‘unrestricted’ strategy, we allow for both (1) longing the futures and shorting the spot and (2) shorting the futures and longing the spot while in the ‘long-spot only’ strategy, we only allow




Figure F.10: Annualized interest rate from Aave

This figure presents the daily supply and borrowing rate from Aave. We consider three stablecoins: USDT, USDC, and DAI. Each day, we take the average interest rate of the 3 crypto stablecoins to get the proxy for the risk-free funding rate of the arbitrageurs. The averaging removes the idiosyncratic noise in borrowing and supply in individual stablecoins. The sample period is from 2020-01-08 to 2022-11-13.

for shorting the futures and longing the spot. The reason we consider the ‘long-spot only’ strategy is the infrastructure for shorting the spot is not well-developed. So we want to examine the performance of the trading strategy in the presence of such limits to arbitrage.

We find the arbitrage trading strategy has a high Sharpe ratio under high trading costs. The annualized Sharpe ratio for Bitcoin is 1.78 in our sample. They are even higher for other cryptos. The high Sharpe ratio of the trading strategy corroborates our theoretical results that when the price deviation is large enough, the trading would be a random-maturity arbitrage opportunity.

In the year 2022, the deviation between the futures and the spot becomes smaller and less volatile. There seems to be a structural break. We indeed find the trading strategy takes less active positions and significantly lower annualized returns. However, when there is a large enough deviation, the strategy can still generate sizeable Sharpe ratios in trading.

The conclusion and results remain similar if we consider ‘long-spot only’ trading strategies where only shorting the futures and longing the spot is allowed. Considering this one-sided trade slightly lowers the Sharpe ratio but the algorithm automatically adjusts by increasing the proportion of times being active. The resulting annualized returns increase. In the year 2022, since most of the time, the futures price is below the spot and we don’t allow




Figure F.11: Trading Strategy Visualization: Random-maturity Arbitrage

This figure presents the random-maturity arbitrage strategy motivated by the theory. The orange, green, and purple lines correspond to the open-position threshold under high, medium, and low trading costs. The red line represents the close-position threshold.

longing the futures and shorting the spot trade, the proportion of time the strategy is active is very low.

Table F.9 further reports the performance of the trading strategy under

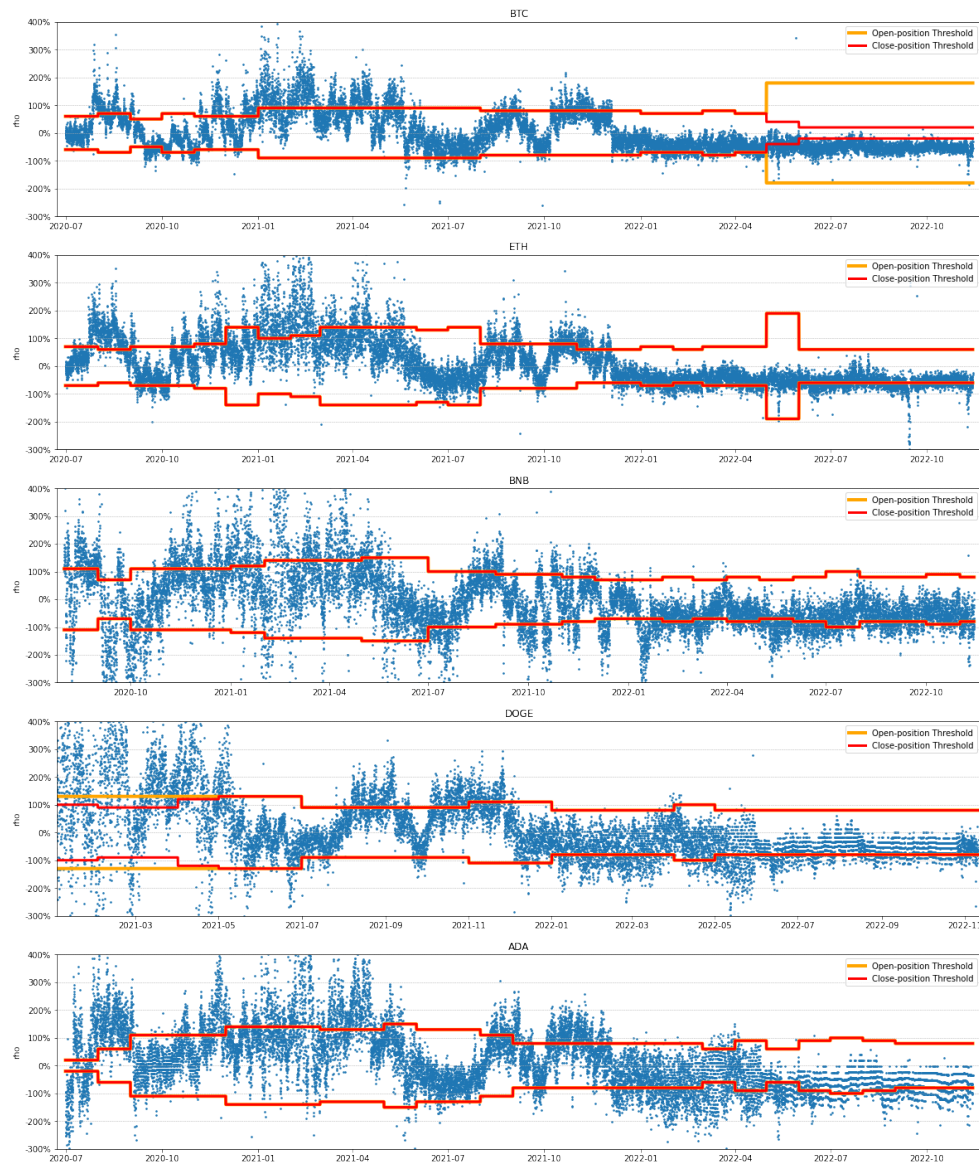


Figure F.12: Trading Strategy Visualization: No Trading Costs

This figure presents the real-time trading thresholds under no trading costs. The orange line is the open-position threshold and the red line is the close-position threshold. The trading thresholds are determined based on the adjusted SR from the past 6 months.

different trading costs specifications laid out in Table F.2. The results from Table F.8 correspond to the last column in Table F.9. As trading costs increase, our trading strategy dynamically adjusts by lowering the proportion

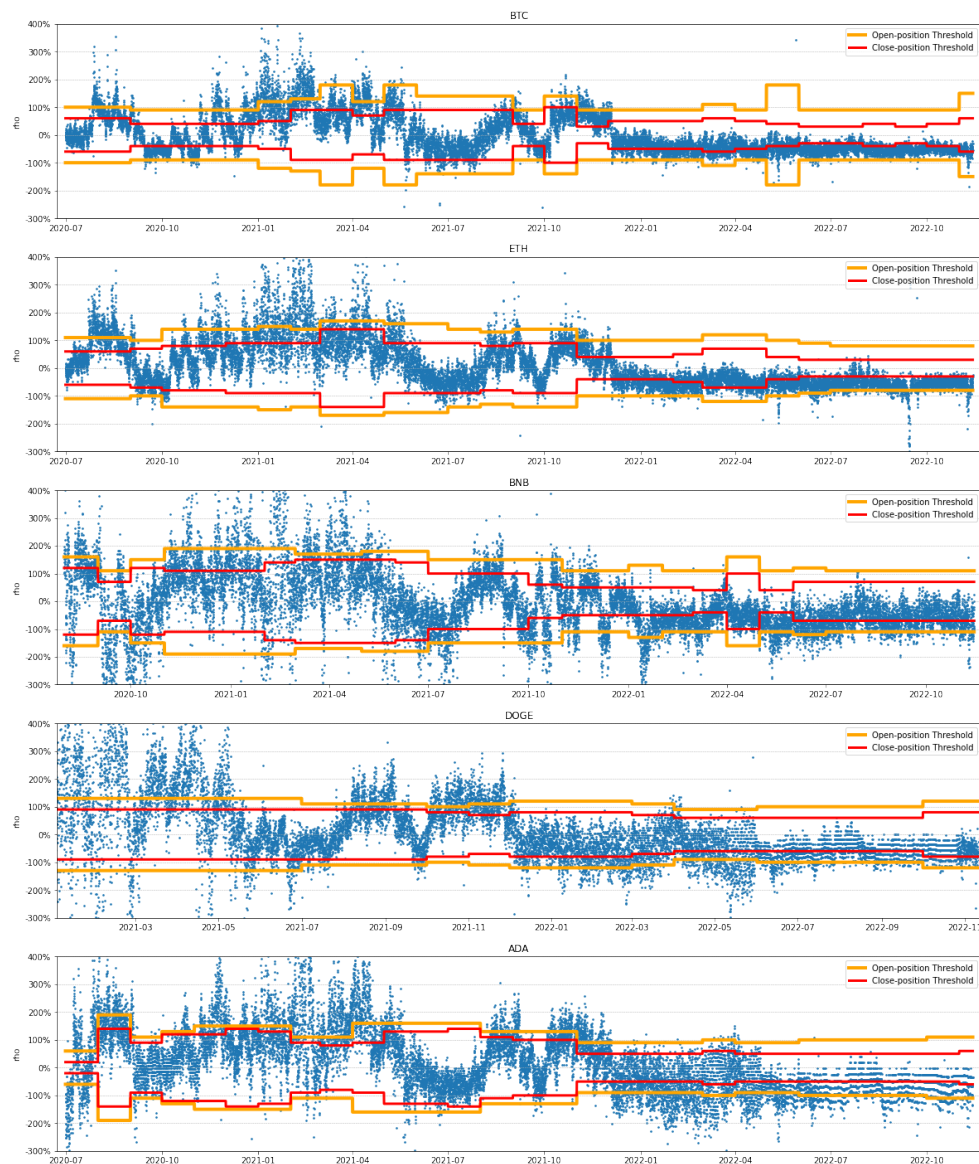


Figure F.13: Trading Strategy Visualization: Low Trading Costs

This figure presents the real-time trading thresholds under low trading costs (2.25 bps for spot trading and 0.18 bp for futures trading). The orange line is the open-position threshold and the red line is the close-position threshold. The trading thresholds are determined based on the adjusted SR from the past 6 months.

of time being active. The annualized return decreases and annualized standard deviation are of similar magnitude across different trading costs. As a result, SR decreases as trading costs increase. Under no trading cost, we see a

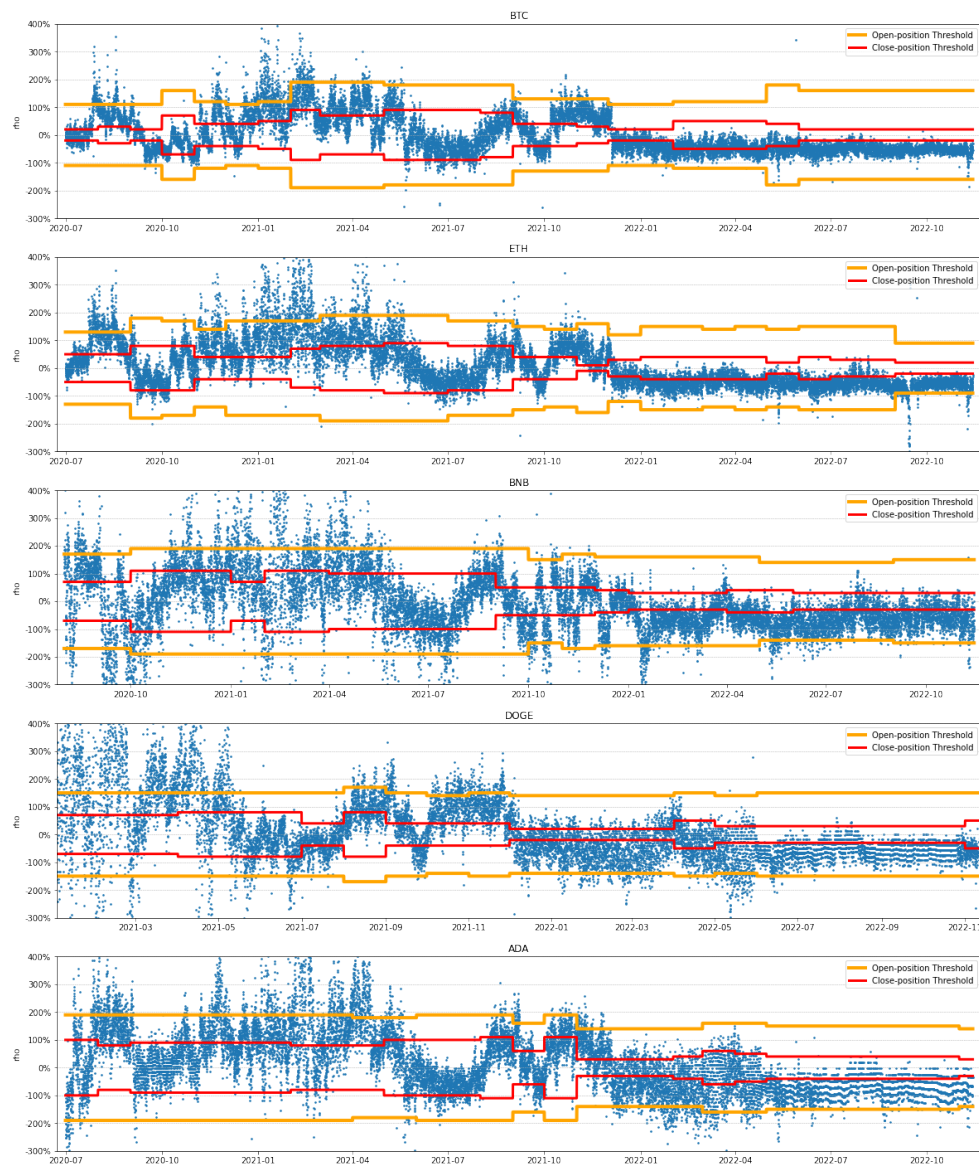


Figure F.14: Trading Strategy Visualization: Medium Trading Costs

This figure presents the real-time trading thresholds under medium trading costs (4.5 bps for spot trading and 0.72 bp for futures trading). The orange line is the open-position threshold and the red line is the close-position threshold. The trading thresholds are determined based on the adjusted SR from the past 6 months.

Sharpe ratio of 6.72 for BTC and Sharpe ratios above 10 for all other cryptos. This also confirms our theoretical results that when there is no trading cost, any deviation of perpetual price from the no-arbitrage benchmark would be a

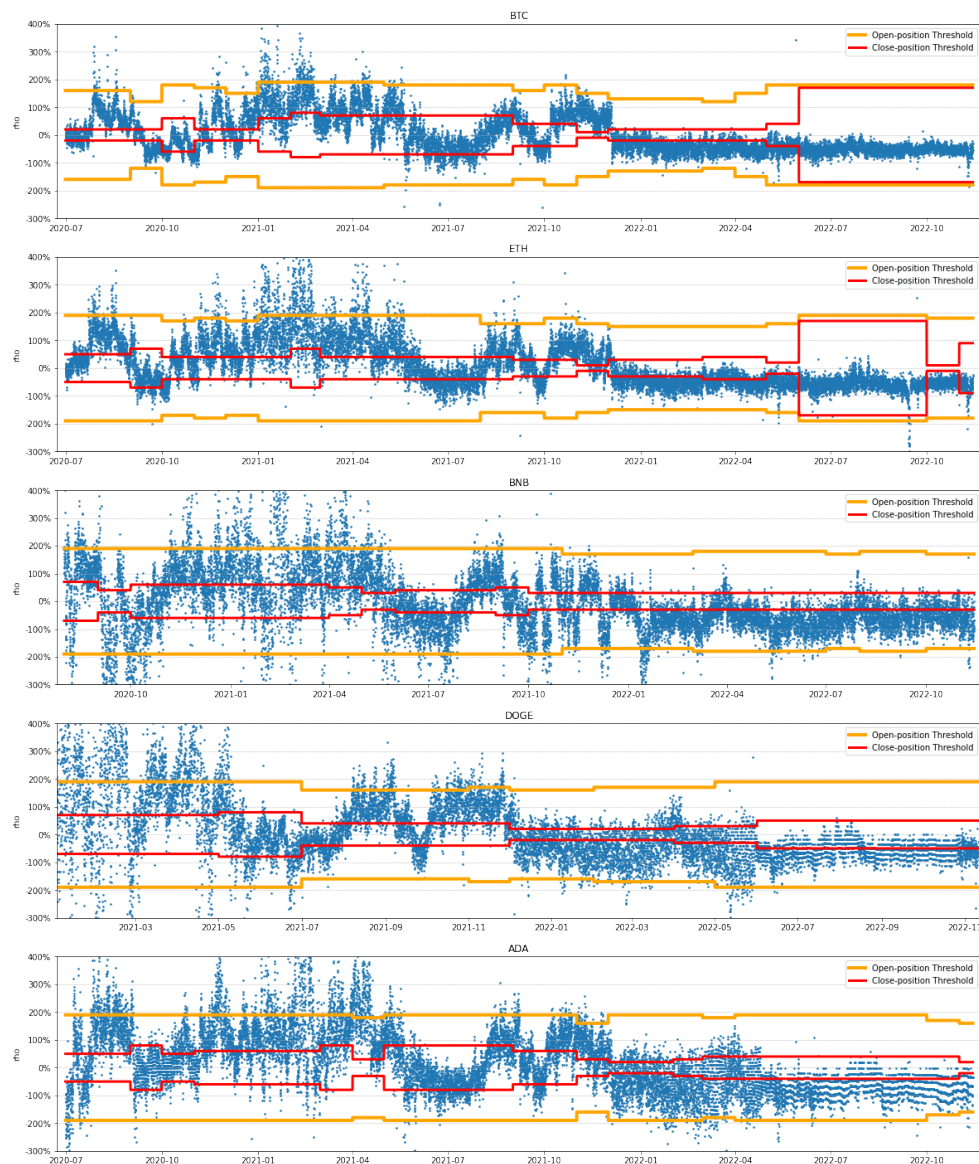


Figure F.15: Trading Strategy Visualization: High Trading Costs

This figure presents the real-time trading thresholds under high trading costs (6.75 bps for spot trading and 1.44 bps for futures trading). The orange line is the open-position threshold and the red line is the close-position threshold. The trading thresholds are determined based on the adjusted SR from the past 6 months.

random-maturity arbitrage opportunity.

Comparing Table F.9 with Table F.4, we find under low and no trading costs, the potential incremental benefits from using a data-driven two-




Figure F.16: Two-threshold trading strategy: Bitcoin, high trading costs

This figure presents real-time trading thresholds of the arbitrage strategy we implement for BTC under high trading costs. Each blue dot in the figure represents the annualized deviation of futures from the spot $\rho = \kappa(f - s) - r$. The orange line is the open-position threshold and the red line is the close-position threshold.

threshold trading rule are much larger compared to higher trading costs. The intuition can be easily illustrated by comparing the empirical real-time trading thresholds in the theory-motivated strategy and the data-driven strategy under low trading costs (Figure F.11 and Figure F.12). With persistent ρ the data-driven approach can find the optimal close-position threshold corresponding to the level of ρ while the theory-motivated one is too conservative in setting a close-position threshold equal to 0.

There are two sources of trading profit for our trading strategy: (1) the price convergence; (2) the funding rate payment. In Table F.10, we decompose the return to our trading strategy without trading cost into the two sources. We find across different cryptos, the return from price convergence accounts for a larger proportion of the total profit. For BTC, the price convergence accounts for more than 2/3 of the profits, and for ETH, it accounts for about 3/4. Additionally, the decomposition for the year 2022 generates different patterns, the return due to funding rate plays a much smaller role across different cryptos. This is because the deviation from perpetual to spot is much smaller during the year. The results suggest when more sophisticated arbitrageurs enter the market, the large deviation of perpetual from the spot would be more of an off-equilibrium outcome. As a result, the funding rate payment will also be small.

In all, our analysis of the trading strategy demonstrates the profitability of perpetual-spot arbitrage trade. The deviation of the futures to spot can

Table F.8 Portfolio Performance: High Trading Cost

		Unrestricted				Long-spot Only			
		2020	2021	2022	All	2020	2021	2022	All
BTC	N	4,416	8,760	7,536	20,712	7,344	8,760	7,609	23,713
	Active %	26.11	18.06	0.29	13.31	30.27	52.45	0.03	28.76
	Return	5.78	8.51	0.13	4.88	6.24	15.21	0.10	7.58
	Volatility	3.74	3.26	0.19	2.74	3.47	6.50	0.18	4.40
	SR	1.55	2.61	0.70	1.78	1.80	2.34	0.54	1.72
ETH	N	4,416	8,760	7,536	20,712	5,880	8,760	7,609	22,249
	Active %	22.98	23.90	12.08	19.40	44.12	54.35	0.09	33.09
	Return	9.23	13.25	0.60	7.79	11.56	18.07	0.61	10.38
	Volatility	2.55	3.73	0.93	2.76	4.29	6.13	0.71	4.45
	SR	3.62	3.55	0.64	2.83	2.70	2.95	0.87	2.33
BNB	N	3,672	8,760	7,536	19,968	3,672	8,760	7,609	20,041
	Active %	41.04	32.32	10.51	25.69	57.30	48.11	1.10	31.94
	Return	24.61	25.33	3.33	16.89	12.23	20.07	0.10	11.05
	Volatility	4.24	4.23	1.32	3.44	3.77	4.80	0.30	3.57
	SR	5.80	5.99	2.52	4.91	3.24	4.18	0.33	3.10
DOGE	N		8,760	7,536	16,296		8,760	7,609	16,369
	Active %		41.66	6.62	25.45		54.26	1.22	29.60
	Return		47.57	0.98	26.02		32.74	0.08	17.56
	Volatility		7.93	0.87	5.85		7.39	0.40	5.42
	SR		6.00	1.14	4.45		4.43	0.21	3.24
ADA	N	4,416	8,760	7,536	20,712	4,416	8,760	7,609	20,785
	Active %	34.24	26.58	7.43	21.24	62.09	55.89	0.00	37.10
	Return	14.15	19.45	0.61	11.46	15.15	18.77	0.00	11.08
	Volatility	3.62	5.13	0.89	3.77	5.07	7.20	0.00	5.23
	SR	3.91	3.79	0.68	3.04	2.99	2.61	0.00	2.12

This table presents the annual return, standard deviations, and Sharpe ratios of the two-threshold trading strategies for the 5 different cryptocurrencies with high trading costs. We also report the proportion of time the trading strategy has an open position. We break down returns into each year and also provide summary stat across all time. The left panel shows the performance of the unrestricted trading strategy where both long and short spot is allowed. The right panel shows the performance of the long-spot-only trading strategy where shorting the spot is not allowed.

also serve as an important measure for the frictions, trading costs, and limits to arbitrage in the market.

Table F.9 Portfolio Performance under Different Trading Costs

		Trading costs			
		None	Low	Medium	High
BTC	Active %	23.20	20.79	16.29	13.31
	Return	15.98	7.42	6.02	4.88
	Volatility	2.38	2.63	2.61	2.74
	SR	6.72	2.82	2.31	1.78
ETH	Active %	32.34	25.85	20.84	19.40
	Return	23.48	11.80	9.41	7.79
	Volatility	2.27	2.27	2.69	2.76
	SR	10.32	5.20	3.50	2.83
BNB	Active %	37.76	30.90	25.52	25.69
	Return	54.50	30.40	22.07	16.89
	Volatility	3.26	3.02	3.12	3.44
	SR	16.72	10.07	7.07	4.91
DOGE	Active %	36.62	39.81	32.46	25.45
	Return	72.95	46.87	34.45	26.02
	Volatility	5.65	5.60	5.81	5.85
	SR	12.90	8.37	5.92	4.45
ADA	Active %	40.00	43.22	24.38	21.24
	Return	53.06	29.63	16.53	11.46
	Volatility	3.22	3.33	3.26	3.77
	SR	16.47	8.90	5.07	3.04

This table presents the portfolio performance under different trading cost specifications. The fee for spot is 0, 1.5 bps, 3.75 bps, and 5.25 bps for the 4 trading costs specifications. The fee for the futures is 0, 0, 0.54 bp, 1.08 bps for the 4 trading costs specifications.

C. High-frequency Event Study of Funding Rate Payment

In this part, we provide a high-frequency event study around the funding rate payment time. We hope to further understand the microstructure effect of the funding rate mechanism on perpetual futures pricing. We obtain 1-min BTC, ETH, and DOGE price data from Kaiko. We then calculate the return from a strategy that taps into the funding rate payment with a window of -4 hours and +4 hours around the funding rate payment. When the funding rate is positive, we stay on the short side of the perpetual futures to receive the funding rate and long the spot to hedge the risk. When the funding rate is negative, we do the trade in the opposite direction.

Figure F.17 presents the average ex-funding-rate cumulative returns of

Table F.10 Return Decomposition: Price Convergence vs Funding Rate Payment

		2020	2021	2022	All
BTC	Return	17.70	25.69	3.69	15.98
	Price	11.77	16.10	3.58	10.62
	Funding	5.93	9.59	0.11	5.36
ETH	Return	22.51	30.99	15.31	23.48
	Price	14.22	21.86	14.36	17.50
	Funding	8.29	9.13	0.95	5.98
BNB	Return	74.35	64.93	32.71	54.50
	Price	61.16	56.63	28.70	46.92
	Funding	13.19	8.29	4.01	7.58
DOGE	Return		96.17	45.95	72.95
	Price		82.50	45.58	65.43
	Funding		13.68	0.36	7.52
ADA	Return	44.24	56.21	54.57	53.06
	Price	31.81	45.69	53.12	45.44
	Funding	12.43	10.52	1.44	7.62

This table decomposes the portfolio return into the part due to price convergence and the part due to funding rate payment.

such a strategy around the time of funding rate payment for BTC, ETH, and DOGE. We see when the funding rate is positive, there is a significant price drop. This reflects the efficiency of the market. When traders observe almost certain positive funding rate payments. They have incentives to short the futures and long the spot. This behavior results in a large decline in spot-perpetual return, which offsets the profits from gaining the funding rate.

On the other hand, we observe less-significant price drops around negative funding rate payments. This can be due to 2 reasons: (1) in our sample, only around 15% of the sample has a negative funding rate and the magnitude is smaller than the positive funding rate. A small sample would result in noisy estimates making it more difficult to discover the patterns in the positive case. (2) To implement the arbitrage strategy, the traders need to long the futures and short the spot. However, the infrastructure for shorting the spot is not well-developed for the crypto spot market. Trading in the futures market alone would expose the traders to too much risk. Because of these limits to arbitrage, we observe a more muted price decline around negative funding rate payments.

Even if the price shows a large decline at high-frequency, it is worth inspecting how much trading profit is left from exploiting the funding rate payment. We implement a high-frequency funding rate arbitrage strategy. Five minutes

before the funding rate payment, we are going to open the position in the futures and spot market. We close the position 5 minutes after the funding rate payment.

We found that the average one-time return from implementing the strategy is about 1.5 bps, which sums up to 16.4% per year. However, implementing the strategy requires round-trip trading in the spot market which has a total trading cost of 3 bps even under low trading cost specification. Therefore after accounting for trading costs, the trading profits will be attenuated. The price movement in perpetual and spot attenuates about $\frac{1}{3}$ of the original profits from receiving the funding rate across the three different cryptos.




Figure F.17: High-frequency Event Study around the Funding Rate Payment time for BTC, ETH, and DOGE

This figure shows cumulative returns to strategies that try to exploit the funding rate payment at the 1-minute frequency for BTC, ETH, and DOGE. The strategies returns reported in the figure do not include the funding rate payment, similar to the ex-dividend returns to stock around the dividend payment. The blue line shows cumulative return to a strategy that shorts the futures and longs the spot when the funding rate is positive. The orange line shows cumulative return to a strategy that longs the futures and shorts the spot when the funding rate is negative.

Paper References

- [1] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core. *Whitepaper*, 2021. URL <https://uniswap.org/whitepaper-v3.pdf>.
- [2] Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, 2010.
- [3] Carol Alexander, Jaehyuk Choi, Heungju Park, and Sungbin Sohn. Bitmex bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1):23–43, 2020.
- [4] Dan Amiram, Evgeny Lyandres, and Daniel Rabetti. Cooking the order books: Information manipulation and competition among crypto exchanges. *SSRN Electronic Journal*, 2021. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3745617.
- [5] Guillermo Angeris, Tarun Chitra, Alex Evans, and Matthew Lorig. A primer on perpetuals. *arXiv*, 2022. URL <https://arxiv.org/abs/2209.03307>.
- [6] Markus K Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *The review of Financial Studies*, 22(6):2201–2238, 2009.
- [7] John Cochrane. *Asset Pricing: Revised Edition*. Princeton University Press, 2009.
- [8] Lin William Cong, Zhiheng He, and Ke Tang. Staking, token pricing, and crypto carry. *SSRN Electronic Journal*, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4059460.
- [9] Lin William Cong, George Andrew Karolyi, Ke Tang, and Weiyi Zhao. Value premium, network adoption, and factor pricing of crypto assets. *SSRN Electronic Journal*, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3985631.
- [10] Lin William Cong, Xi Li, Ke Tang, and Yang Yang. Crypto wash trading. Technical report, National Bureau of Economic Research, 2022.

- [11] Riccardo De Blasis and Alexander Webb. Arbitrage, contract design, and market structure in bitcoin futures markets. *Journal of Futures Markets*, 42(3):492–524, 2022.
- [12] Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from Covered Interest Rate Parity. *Journal of Finance*, 73(3):915–957, 2018.
- [13] Darrell Duffie. Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance*, 65(4):1237–1267, 2010.
- [14] Alex Ferko, Amani Moin, Esen Onur, and Michael Penick. Who trades bitcoin futures and why? *Global Finance Journal*, 2022.
- [15] Nicolae Garleanu and Lasse Heje Pedersen. Margin-based asset pricing and deviations from the law of one price. *The Review of Financial Studies*, 24(6):1980–2022, 2011.
- [16] Denis Gromb and Dimitri Vayanos. The dynamics of financially constrained arbitrage. *The Journal of Finance*, 73(4):1713–1750, 2018.
- [17] Eyal Hertzog, Guy Benartzi, Galia Benartzi, and Omri Ross. Bancor protocol, continuous liquidity for cryptographic tokens through their smart contracts. *Whitepaper*, 2017. URL https://2017.iswi.org/wp-content/uploads/sites/17/2017/05/Bancor_protocol-MAY5.pdf.
- [18] Ralph SJ Koijen, Tobias J Moskowitz, Lasse Heje Pedersen, and Evert B Vrugt. Carry. *Journal of Financial Economics*, 127(2):197–225, 2018.
- [19] Yukun Liu and Aleh Tsyvinski. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6):2689–2727, 2021.
- [20] Yukun Liu, Aleh Tsyvinski, and Xi Wu. Common risk factors in cryptocurrency. *The Journal of Finance*, 77(2):1133–1177, 2022.
- [21] David O Lucca and Emanuel Moench. The pre-fomc announcement drift. *The Journal of Finance*, 70(1):329–371, 2015.
- [22] Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319, 2020.
- [23] Lasse Heje Pedersen. *Efficiently Inefficient: how smart Money invests and Market Prices are determined*. Princeton University Press, 2019.
- [24] Maik Schmeling, Andreas Schrimpf, and Karamfil Todorov. Crypto carry. *SSRN Electronic Journal*, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4268371.

- [25] Robert J Shiller. Measuring asset values for cash settlement in derivative markets: hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931, 1993.
- [26] Artem Streltsov and Qihong Ruan. Perpetual price discovery and crypto market quality. *SSRN Electronic Journal*, 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4218907.