



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details



Topics in Crypto Asset and Blockchain Finance

Michael Dakos
University of Sussex

A thesis submitted for the degree of

Doctor of Philosophy

August 2022

Acknowledgements

During my PhD the most common reading material included academic papers and textbooks, but as I approach the end of this journey, it is a poem by C. P. Cavafy – *Ithaca*, that comes to mind:

*As you set out for Ithaca
hope that your journey is a long one,
full of adventure, full of discovery.*

*Laistrygonians and Cyclops,
wild Poseidon – you won't encounter them
unless you bring them along inside your soul,
unless your soul sets them up in front of you.*

*Hope that your journey is a long one.
May there be many summer mornings when,
with what pleasure, what joy,
you enter harbours seen for the first time,
to learn and go on learning from their scholars.*

*Arriving there is what you're destined for.
But don't hurry the journey at all.
Better if it lasts for many years.*

*Wise as you will have become, so full of experience,
you'll have understood by then what these Ithacas mean.*

I wish to express my sincere thanks and gratitude to my main supervisor Carol Alexander, for her patient guidance and kind support throughout this journey. Working alongside Carol has been a privilege and an honour, and has helped me immensely in growing both professionally and personally.

Many thanks also go to Martin White, for the consistent provision of helpful comments together with his technical expertise and insights.

Last but not least, I want to thank my family, friends and fellow PhD students for sharing this journey with me.

Declaration of original authorship

I hereby declare that this thesis is my own work and that no other sources have been used except those clearly indicated and referenced. This thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Athens, 1 August 2022
Michael Dakos

Abstract

This thesis contributes to the crypto asset and blockchain empirical finance literature in three key areas: (i) market risk modelling, by developing simple volatility models which exhibit equal forecasting ability in terms of crypto asset tail risk measure and volatility forecasts, when compared against complex models; (ii) market manipulation, by extending a methodology derived from securities fraud litigation studies to identify blockchain transactions with a manipulative effect on crypto asset prices; (iii) crowdfunding via token offerings, by identifying factors of fundraising success using regression models, and exploring how these factors vary across time. Each of the above contributions is developed in a separate chapter.

Firstly, the market risk modelling chapter provides extensive backtests of hourly and daily Value-at-Risk and Expected Shortfall forecasts regarded as best practice in the industry and used for regulatory approval. Results demonstrate that simpler models in the EWMA class are just as accurate as GARCH models for VaR and ES forecasting, and similarly when using average scores generated from proper univariate and multivariate scoring rules.

Secondly, the market manipulation chapter examines large blockchain transactions of the tether stablecoin and assesses whether they produce positive abnormal returns for bitcoin prices. The methodology is adapted from single-firm event studies used in securities fraud litigation, using regression factor models. The chapter's findings can be useful in determining materiality and estimating damages in legal cases of crypto asset market manipulation.

Finally, the tokenomics of crowdfunding chapter examines the fundraising success of token offerings for the 2017 – early 2022 period, constituting one of the most comprehensive studies in this topic. We proxy fundraising success with the amount of funding raised and also by minimum funding target exceedance. Success factors are derived from the venture, token and offering characteristics, as well as additional common factors such as the price level of ether and the launchpad platforms used. The findings of this chapter provide insights as to the evolution of token offering success factors, with the choice of launchpad platform emerging as a new and significant factor and to some extent overshadowing the determinants previously documented in the relevant literature.

Contents

1	Introduction	1
2	Crypto Asset Market Risk Modelling	20
2.1	Introduction	21
2.2	Literature Review	23
2.2.1	Models	27
2.2.2	In-Sample Performance	31
2.2.3	Out-of-Sample Performance	33
2.2.4	Focused Discussion and Contributions	38
2.3	Methodology	42
2.3.1	Volatility and Covariance Models	43
2.3.2	Value-at-Risk and Expected Shortfall	46
2.3.3	Distribution Forecast Evaluation	52
2.4	Data	55
2.5	Empirical Results	59
2.5.1	Daily Forecast Evaluation	59
2.5.2	Hourly Forecast Evaluation	71
2.6	Summary and Conclusions	80
3	Blockchain Transactions and Crypto Asset Market Manipulation	84
3.1	Introduction	85
3.2	Literature Review	87
3.3	Methodology	99
3.4	Data	105
3.4.1	Market Data	106
3.4.2	Blockchain Data	109
3.4.3	Tether on the Omni Layer	116
3.5	Empirical Results	126
3.6	Summary and Conclusions	130
4	The Tokenomics of Crowdfunding	132
4.1	Introduction	133
4.2	Literature Review & Hypotheses Development	136
4.2.1	Literature Review	136
4.2.2	Hypotheses Development	142
4.3	Data	150
4.3.1	Variables	152

4.3.2	Sample Statistics	158
4.4	Models	161
4.5	Empirical Results	165
4.5.1	Sample Period 2017 – 2020	165
4.5.2	Sample Period 2021 – 2022	177
4.6	Summary and Conclusions	184
5	Summary and Conclusions	189
A	Supplementary Material – Crypto Asset Market Risk Modelling	A1
B	Supplementary Material – Blockchain Transactions and Crypto Asset Market Manipulation	A4
C	Supplementary Material – The Tokenomics of Crowdfunding	A8

List of Figures

1.1	Stablecoin market capitalization	8
1.2	Crypto asset market capitalization	10
1.3	Bitcoin (BTC/USD) price	11
1.4	Ether (ETH/USD) price	11
1.5	Total value locked in DeFi platforms	12
2.1	Daily log returns of bitcoin, ether, ripple and litecoin	56
2.2	Hourly log returns of bitcoin, ether, ripple and litecoin	58
3.1	Hourly BTC/USD price and USDT supply	105
3.2	Hourly returns	108
3.3	Distribution of USDT outflows from the tether treasury	114
3.4	USDT outflows from the tether treasury to exchanges	115
3.5	USDT outflows from the tether treasury to unknown entities	115
3.6	Omni Layer data retrieval execution times	120
4.1	Token offerings and funding raised between 2017 – 2022	153
4.2	Token offerings per launchpad platform and coefficients	183
4.3	Token offerings per category and coefficients	184

List of Tables

1.1	Overview of the crypto asset and blockchain finance academic literature	14
1.2	Literature streams in crypto asset and blockchain finance	16
2.1	Literature on crypto asset volatility and covariance forecasting	23
2.2	Sample statistics of daily returns on bitcoin, ether, ripple and litecoin	57
2.3	Sample statistics of hourly returns on bitcoin, ether, ripple and litecoin	59
2.4	Backtesting results for 1-day-ahead left-tail VaR	63
2.5	Backtesting results for 1-day-ahead right-tail VaR	64
2.6	Backtesting results for 1-day-ahead left-tail ES	66
2.7	Backtesting results for 1-day-ahead right-tail ES	67
2.8	Average CRPS of 1-day-ahead univariate density forecasts	70
2.9	Average multivariate scores for 1-day-ahead joint density forecasts . .	72
2.10	Backtesting results for 1-hour-ahead left-tail VaR	76
2.11	Backtesting results for 1-hour-ahead right-tail VaR	77
2.12	Backtesting results for 1-hour-ahead left-tail ES	78
2.13	Backtesting results for 1-hour-ahead right-tail ES	79
2.14	Average CRPS of 1-hour-ahead univariate density forecasts	80
2.15	Average multivariate scores for 1-hour-ahead joint density forecasts .	81
3.1	Literature of the effect of stablecoins on crypto assets	88
3.2	Summary statistics of hourly returns	107
3.3	Correlation matrix of hourly returns	109
3.4	Summary statistics of USDT outflows from the tether treasury	113
3.5	Material USDT outflows from the tether treasury to exchanges	128
3.6	Material USDT outflows from the tether treasury to unknown entities	129
4.1	Empirical results in the token offerings literature	137
4.2	Token offerings sample statistics 2017 – 2020	159
4.3	Token offerings correlation matrix 2017 – 2020	159
4.4	Token offerings (capped) sample statistics 2017 – 2020	160
4.5	Token offerings (capped) correlation matrix 2017 – 2020	160
4.6	Token offerings sample statistics 2021 – 2022	161
4.7	Token offerings correlation matrix 2021 – 2022	161
4.8	Linear regression model parameters 2017 – 2020	167
4.9	Probit regression marginal effects 2017 – 2020	168
4.10	Linear regression parameters for the 2021 – early 2022 period	178

INTRODUCTION

On 3 January 2009, the Sigsript field of the very first [Bitcoin transaction](#) contained the hexadecimal string 5468652054696d65732030332f4a616e2f32303039204368616e63656c6c6f72206f6e206272696e6b206f66207365636f6e64206261696c6f757420666f722062616e6b73, which in ASCII text translates to: ‘The Times 03/Jan/2009 Chancellor on brink of second bailout for banks’, a laconic critique by the creator(s) of Bitcoin on the handling of the 2008 global financial crisis. Bitcoin is referred to (Antonopoulos, 2017) as ‘a network of trust that could also provide the basis for so much more than just currencies’. Since 2009, the technological innovations first introduced by Bitcoin have inspired a wave of disruptive innovation that ‘gives every citizen [...] the ability to innovate in terms of financial instruments, payment systems, and banking’ (Antonopoulos, 2016).

The emergence of crypto assets began in 2008 with the need for ‘an electronic payment system based on cryptographic proof instead of trust’, as highlighted by the pseudonymous author(s) (Nakamoto, 2008) of the Bitcoin whitepaper. Crypto assets have since proliferated at an extremely high rate, largely driven by the low entry barrier due to the open-source nature of most crypto asset software projects, and are by now considered as a distinct asset class (Burniske and Tatar, 2018). The earliest record of the crypto asset information aggregation website [Coinmarketcap](#) in 2013 listed only 7 crypto assets; by comparison, at the time of writing, the ECB financial stability review (Hermans et al., 2022) reports 16,000 crypto assets in existence, Coinmarketcap lists upwards of 9,000, while a more conservative estimate of 3,200 is provided by [Cryptocompare](#), another aggregation website.

Similarly, at the time of writing, crypto asset use cases extend well beyond ‘cryptocurrencies’ such as the value transfer system application introduced by Bitcoin.¹ A crypto asset taxonomy report from Cryptocompare (2018) makes an initial distinction based on fungibility, i.e. by whether units of a crypto asset are mutually interchangeable;² focussing on fungible crypto assets, the taxonomy distinguishes between payment, utility and asset-security crypto assets, extending a similar framework introduced first by FINMA (2018), the Swiss financial market supervisory authority. The assumed primary functionality of payment crypto assets, such as bitcoin and litecoin, is to be an alternative electronic cash system, also potentially serving as a store of value. Utility crypto assets provide digital access to an application or service, such as ether, EOS and dot which are the native ‘tokens’ used to access and utilize respectively the Ethereum, EOS and Polkadot decentralized platforms. Asset-security crypto assets are analogous to equities, bonds or financial derivatives in that they encompass a participation in real physical underlying assets, companies, or earnings streams, and are usually subject to much stricter regulation, contrary to payment and utility crypto assets.

The academic literature on crypto asset and blockchain finance is equally burgeoning, with more than 2,200 relevant articles published since 2013 in finance, economics, econometrics, business, management and accounting journals, and many more conference and working papers.

This doctoral thesis develops three distinct topics in the area of crypto asset and blockchain empirical finance: (i) crypto asset market risk modelling; (ii) crypto asset market manipulation; and (iii) crowdfunding via token offerings. The remainder of this introductory chapter provides brief definitions and an introduction to blockchain, crypto assets and crypto asset markets, a broad overview of the literature on crypto

¹Varied terminology is used by practitioners and academics, such as the earlier term ‘cryptocurrency’, sometimes shortened to ‘crypto’, as well as the more general term ‘crypto asset’ which is preferred throughout this thesis. Note also, as clarified by Burniske and Tatar (2018), that Bitcoin – capital ‘B’, denotes the platform that carries upon it the programmable money known as bitcoin – lowercase ‘b’.

²Note that non-fungible crypto assets which bear characteristics of uniqueness are not covered in this thesis, but nonetheless present interesting use cases such as digital collectibles and art.

asset and blockchain finance and finally discusses the contributions of this thesis to the relevant literature.

Blockchain and Crypto Assets

The concept of a blockchain traces back to the Bitcoin whitepaper (Nakamoto, 2008) which refers to a ‘chain of hash-based proof-of-work’. Blockchain is a distributed ledger of timestamped transactions, where data are stored in blocks which contain information on one or more transactions; it stores data in an immutable manner and its contents are auditable and verifiable by some or all stakeholders in the network. For instance, the Bitcoin blockchain is the ledger that includes the entire history of bitcoin transactions and circumvents the double-spend problem, i.e. the risk that a unit of bitcoin could be spent more than once due to its lack of a physical substance that would render such an attempt impossible.

The discussion provided in this section focuses primarily on public blockchains which are used by the overwhelming majority of crypto assets. However, before expanding on public blockchains, it is important to also note the existence of private or permissioned blockchains and distributed ledger technologies (DLT); these are not covered in this thesis, but have nonetheless attracted significant attention in the space of enterprise software solutions and also in financial applications such as central bank digital currencies (CBDCs). This is partially driven by potential drawbacks of public blockchain systems such as reduced privacy due to the public nature of all transactions, and scalability issues due to the data storage and processing requirements.

Private blockchains and DLTs address these issues (Platt, 2017) by ‘relaxing’ the assumptions on transparency, security and immutability, and imposing data segregation rules – i.e. limiting transaction visibility into channels or sub-ledgers, and also channel independence – i.e. only requiring consensus on the state of a channel/sub-ledger. Prominent private blockchains and relevant applications include

[Corda](#) by the R3 enterprise consortium, [Quorum](#) by ConsenSys, [Hyperledger](#) by the Linux Foundation and [Azure](#) by Microsoft.

The suggestion of a public blockchain structure by Nakamoto (2008) provides a solution to the ‘byzantine generals’ problem, a common problem in computer networks involving the transfer of sensitive information over an unreliable network. The manner by which participants in such a network reach an agreement as to the current state of the blockchain ledger is called a consensus algorithm. Several approaches exist to this issue, such as the proof-of-work approach used by Bitcoin and several other crypto assets, the proof-of-stake approach, and other hybrid approaches that include delegated proof-of-stake (DPoS), proof-of-burn, proof-of-authority etc.

In proof-of-work consensus, the users – called miners, create blocks of newly-executed transactions and then expend energy and computational processing power in order to solve an asymmetric difficulty problem by trial-and-error (Antonopoulos, 2017).³ The first miner to create a new block of transactions is rewarded with newly-created units of the crypto asset, while the efforts of all other miners are wasted. The difficulty of the proof-of-work problem is set arbitrarily so that the average solution time is 10 minutes, which renders any attempt to alter previous blocks in the blockchain very hard computationally, if not impossible.

[Proof-of-stake](#) is a different solution to the byzantine generals problem in which holders of the crypto asset are perceived as holding stakes in the network, and may cast stake-weighted votes to validate and include blocks of new transactions in the blockchain. Users – called validators, stake their crypto asset holdings as collateral, to be forfeited in case of laziness – i.e. creating empty blocks with no transactions, or dishonesty – i.e. including invalid or double-spend transactions. Contrary to proof-of-work, proof-of-stake does not involve wasted effort and is perceived on the whole to be less energy-intensive. Proof-of-stake is a more complex consensus mechanism than proof-of-work, and is considered to be more secure and better for

³The asymmetry of the proof-of-work problem is that the mathematical/cryptographic problem is computationally hard and time-consuming to solve, but the validity of any potential solution can be checked quickly and easily.

implementing new scaling solutions. Ethereum began with a proof-of-work consensus in 2014 and has been developing a proof-of-stake implementation for several years, but has not fully transitioned to it at the time of writing.

A crypto asset is a digital asset that uses web-based distributed public ledgers to establish proof-of-ownership. Crypto assets may be deployed on their own native public blockchain, or they may ‘piggy-back’ on other already-established blockchains as ‘tokens’.⁴ Regardless of the native-blockchain vs. token distinction, it is also important to note that the total supply of any crypto asset and also the method(s) of new supply issuance can be set arbitrarily by the crypto asset developers, in stark contrast to *fiat* currencies and more akin e.g. to equities.

For instance, new supply of bitcoin (BTC) is provided as a reward for mining; it started at 50 BTC per block in 2009 and is set to halve approximately every four years, causing the total supply of bitcoin to be capped at approximately 21 million BTC. By comparison, at the launch of Ethereum in 2014, a total of 72 million ether (ETH) – the Ethereum platform’s native crypto asset, were ‘pre-mined’ and sold to prospective investors, while newly-created units of ether – approximately 2 ETH per new block created, are awarded to blockchain validators. Moreover, and contrary to bitcoin, there is theoretically no cap on the total amount of ether units that can be created. However, in 2021 a potentially deflationary element was added to ether’s supply schedule via ‘burning’ – i.e. a destruction mechanism, of ether transaction fees to discourage Ethereum miners from raising the minimum Ethereum transaction fee for users.

The specialised scope of blockchain in the context of Bitcoin, as a ledger of value transfer transactions, is generalised by the Ethereum platform into a cryptographically secure, transaction-based state machine. Ethereum was launched in 2014 and is described in the corresponding yellow paper (Wood, 2014) as a project which ‘attempts to build the generalised technology on which all transaction-based state machine

⁴Native-blockchain crypto assets are often referred to as ‘coins’ – an extension of the ‘cryptocurrency’ terminology; however, in this thesis we avoid the term ‘coins’, as it alludes primarily to the payment crypto asset use case, while native-blockchain crypto assets may also belong to the utility or asset-security type.

concepts may be built'. In other words, Ethereum and similar platforms can execute computer code in an immutable and publicly verifiable manner, in the form of smart contracts which were first introduced theoretically by Szabo (1997); smart contracts can be deployed and executed by every participant in the network on the Ethereum virtual machine (EVM), while all smart contract and ether transactions are embedded in the Ethereum blockchain in return for fees payable to blockchain validators in the form of ether.

Smart contracts can include crypto asset storage in escrow, payments in very complex conditional or repetitive structures and even the creation of 'tokens', i.e. crypto assets deployed on a non-native blockchain, potentially resulting in financial contracts or even in decentralized corporate and organizational structures. The generalised nature of smart contract platforms such as Ethereum have given rise to numerous crypto asset and blockchain applications, which include – but are by no means limited to: physical asset and commodity tokenization, content distribution platforms, data exchange platforms, data processing services, decentralized crypto asset exchanges, decentralized domain name systems, decentralized marketplaces, decentralized social media platforms, decentralized applications, digital advertising, digital asset management, platforms of digital content production, distribution and management, distributed computing, distributed crowdfunding, distributed data storage, distributed virtual private networks, gaming and gambling platforms, identity verification applications, in-game currency applications, instant messaging, internet of things applications, loyalty and reward points marketplaces, medical record storage networks, peer-reviewed knowledge platforms, prediction markets, price-pegged crypto assets such as stablecoins, and privacy-based payments systems.

Of the above applications, it is important to expand on stablecoins and decentralized crowdfunding via token offerings, which both play an important role in crypto asset markets and are the core focus of chapters 3 and 4.


Stablecoins constitute a subset of crypto assets whose price is stabilised via price-pegs or other means. The main types of stablecoins are asset-backed and

algorithmic. In asset-backed stablecoins, one or more traditional assets – such as US dollar deposits, or crypto assets are held in reserve to support the stablecoin’s price-peg. Asset-backed stablecoins may be over-collateralized or collateralized on a 1-to-1 basis; for instance, DAI holds a dynamic reserve of crypto assets in excess of its total supply to act as a hedge against market volatility; on the contrary, tether (USDT) was supposed to hold USD deposits on a 1-to-1 basis, but has recently disclosed that its collateral also consists of commercial paper to a significant extent. Algorithmic stablecoins are uncollateralized and instead employ a set of rules that attempt to maintain a dynamic supply-demand equilibrium. For instance, the Terra USD (UST) stablecoin has a built-in arbitrage mechanism with its ‘sister’ crypto asset Luna – the native crypto asset of the Terra blockchain; in short, the creation of new UST supply requires the ‘burning’ – i.e. destruction of Luna supply.

Stablecoins play a vital role in the crypto asset ecosystem by providing a liquidity solution in the form of a ‘dollar proxy’ for crypto asset exchanges, many of which are unregulated and don’t accept fiat currency deposits. For instance since 2019, over 70% of bitcoin’s spot exchange daily traded volume is against a stablecoin – tether (USDT), as reported by [Cryptocompare](#). The systemic role of stablecoins in the crypto asset market can also be seen in the total stablecoin supply, which has increased significantly in recent years, as shown in Figure 1.1.

Another recent and striking example of stablecoins’ vital role in the crypto asset ecosystem is the Terra USD/Luna meltdown in May 2022 which caused a crash in the entire crypto asset market. A possible cause of the meltdown is assumed to be a large volume of Terra USD (UST) withdrawals on the Terra ecosystem followed by an immediate sell-off on crypto asset exchanges; this created downward price pressure for the UST ‘stablecoin’. Traders attempted to exploit the UST/Luna arbitrage opportunity, which would have restored the UST dollar peg, but the mechanism collapsed due to a \$100 million daily cap on UST burning for Luna. As UST could not maintain its dollar peg, a massive sell-off ensued and the ‘dollar-pegged’ UST even traded for a few cents. The Terra/Luna developers subsequently attempted to

Figure 1.1: Stablecoin market capitalization



Distribution of stablecoin market capitalization across largest stablecoins in billion USD (upper panel) and relative to total capitalization (lower panel) between January 2017 – June 2022 based on monthly frequency data obtained from Coinmarketcap. The stablecoins included are Tether (USDT), USD Coin (USDC), Binance USD (BUSD), Terra (UST), Dai (DAI), and the Other category further includes TrueUSD (TUSD), Pax Dollar (USDP), Neutrino USD (USDN), Decentralized USD (USDD), Fei USD (FEI) and Gemini Dollar (GUSD).

restore the UST dollar peg via a \$1 billion UST ‘burning’ to decrease the token’s supply, but the attempt was unsuccessful: at the time of writing, UST – rebranded as Terra Classic USD, trades at approximately \$0.03.

Token offerings, i.e. crowdfunding via crypto assets have attracted significant investor interest in 2017 – 2018 and also more recently in 2021 – 2022.⁵ A CoinDesk article (Feign, 2021) defines token offerings as a method of fundraising for early-stage crypto asset projects, where a start-up venture mints a certain quantity of its own crypto asset or token and offers it to prospective investors in exchange for other established crypto assets – such as ether, and sometimes also *fiat* currencies. [Cryptorank](#) – a relevant information aggregation website, lists 3,500 token offerings completed since 2015, while the token offerings research database (Momtaz, 2021)

⁵Various other terminology is used by academics and practitioners to refer to variants of crowdfunding via crypto assets, the most popular of which are initial coin offerings (ICOs), initial exchange offerings (IEOs), initial decentralized exchange offerings (IDOs) etc. Additionally, as both native-blockchain crypto assets and tokens are used in crypto asset crowdfunding, the most general term would be ‘initial crypto asset offering’, but the ICAO acronym is already used by the International Civil Aviation Organization. This thesis therefore follows the more general term ‘token offerings’.

lists 6,400 such ventures. The total funding raised by token offerings since 2015 is approximately \$13 billion in USD-equivalent valuation, based on Cryptorank and other sources. By comparison, a total of \$286 billion was raised via [IPOs in 2021](#) alone, and similarly the valuation of the [global crowdfunding market](#) – excluding token offerings, in 2020 was \$12 billion.

Crypto Asset Markets

The ECB financial stability review of May 2022 (Hermans et al., 2022) indicates increasing investor demand for crypto assets, attributed inter alia to their frequent use as an instrument of speculation, their unique characteristics such as programmability and also potential benefits from portfolio diversification. Crypto asset markets currently represent less than 1% of the global financial system in size, but are nonetheless similar in size to the securitised sub-prime mortgage market that triggered the global financial crisis of 2007-08.

Since 2013, the estimated capitalization of the entire crypto asset market as reported in [Coinmarketcap](#) has increased from \$1 billion in 2013, to almost \$3 trillion in late 2021 and again dropped to \$1 trillion at the time of writing in mid-2022, as also shown in Figure 1.2. By comparison, the current total market capitalization of all constituent stocks in the [Wilshire 5000](#) broad equity market index is approximately \$40 trillion and similarly the [total market capitalization](#) of all companies listed globally has increased from \$65 trillion in 2013 to \$90 trillion in 2022. Moreover, at the time of writing, only approximately 25 crypto assets have a market capitalization comparable with that of a large-cap equity, and the top 10 large-cap crypto assets currently account for more than 80% of the total crypto asset market capitalization.

Crypto assets are considered highly risky and speculative investments (Hermans et al., 2022), with market risk and operational risk identified as the key risk types. Market risk – which is the focus of Chapter 2, is a result of the extremely high volatility exhibited by crypto asset prices – as shown for the dollar price of bitcoin and ether in Figures 1.3 and 1.4. Operational risks include market manipulation

Figure 1.2: Crypto asset market capitalization




Total crypto asset market capitalization in billion USD (regular axis in upper panel, logarithmic axis in lower panel) between July 2010 – July 2022 based on weekly frequency data obtained from Coinmarketcap and [Bitinfocharts](#).

– which is the focus in Chapter 3, and also misleading information, cyber attacks, fraud and scams. In particular, several financial regulators such as the U.S. Securities and Exchanges Commission (SEC) have issued [warnings](#) to potential investors on the risks of participating in token offerings – of which the fundraising success is examined in Chapter 4.

An industry report by Daye et al. (2019) identifies crypto asset exchanges as one of the key participants in crypto asset markets, while a similar blog post by Carter (2018) focusses on centralised spot exchanges (CEXs) and distinguishes between ‘fiat on-ramps’ and ‘altcoin casinos’: fiat on-ramps – such as Coinbase and Kraken, accept fiat currencies and as a result tend to be regulated and comply with know-your-customer and anti-money-laundering (KYC/AML) regulations; in contrast, ‘altcoin casinos’ – such as Binance, don’t accept fiat currencies, primarily offer crypto-to-crypto and crypto-to-stablecoin traded pairs, and tend to be lightly regulated at best.⁶

⁶Additionally, Daye et al. (2019) mention decentralized exchanges (DEXs) and also exchanges offering crypto asset derivative products, which are not covered in this thesis.

Figure 1.3: Bitcoin (BTC/USD) price




Bitcoin BTC/USD price (regular axis in upper panel, logarithmic axis in lower panel) between July 2010 – July 2022 based on daily frequency data obtained from Cryptocompare.

Figure 1.4: Ether (ETH/USD) price



Ether ETH/USD price (regular axis in upper panel, logarithmic axis in lower panel) between August 2015 – July 2022 based on daily frequency data obtained from Cryptocompare.

Figure 1.5: Total value locked in DeFi platforms



Total value locked (TVL) in billion USD – equivalent valuation, in DeFi platforms between November 2018 – June 2022 based on monthly frequency data obtained from [DeFi Llama](#).

Another important participant in crypto asset markets are the so-called ‘DeFi’ (decentralized finance) platforms and applications that emerged in 2020. The DeFi space comprises a variety of smart contract-based financial applications. Common applications include: decentralized exchanges (DEXs) – such as UniSwap, Bancor and Balancer; peer-to-peer lending platforms such as Aave and Compound; prediction markets – such as Augur and Gnosis; and stablecoin issuers – such as Maker (Hertig, 2020). This multitude of financial applications has caused yet another wave of disruption in the financial intermediation industry and allowed end users and retail investors to participate in complex financial transactions that include: yield farming – i.e. lending rate arbitrage, liquidity mining – i.e. the contribution to DEX liquidity pools in exchange for fees, and ‘money legos’ – i.e. the seamless combination of multiple DeFi applications on the same ecosystem to carry out complex investment strategies. DeFi platforms and applications have attracted significant investor interest, as shown in Figure 1.5 which exhibits the evolution of the ‘total value locked’ (TVL) metric – i.e. the dollar-equivalent of all funding onboarded and invested in DeFi applications.

Having examined the key participants in crypto asset markets, we end this overview with a discussion of the latest crypto asset trend – non-fungible tokens (NFTs). While NFTs are in existence since early 2018 with the introduction of Ethereum ERC-721 non-fungible token standard, the relatively recent development of DeFi-related NFT marketplaces together with the emergence of several metaverse-related gaming applications has caused a surge in interest. An NFT represents a full

or fractional ownership deed for digital assets that include – but are not limited to – artwork, audio tracks, photographs and videos, profile pictures, collectibles such as digital trading cards, web domain names, and e-sports, video game and virtual reality items. Additionally, NFTs relate to and extend the concept of the [metaverse](#) – i.e. ‘a future, more immersive version of the internet, which is envisioned to be experienced via avatars in shared 3D spaces’. For instance, ‘The Sandbox’ is an Ethereum-based video game allowing users to own pieces of NFT-based digital land and develop and monetize custom locations and games within the Sandbox environment.

The Crypto Asset and Blockchain Finance Literature

As mentioned previously, the crypto asset and blockchain finance literature is burgeoning, with a total of 2,270 relevant articles published in a total of 534 academic journals.⁷ Of this total, approximately 25% are published in 8 journals as follows: 11% in *Finance Research Letters* and between 2%–3% each in *Economics Letters*, *Research in International Business and Finance*, *International Review of Financial Analysis*, *Applied Economics Letters*, *Financial Innovation*, *North American Journal of Economics and Finance*, and *Technological Forecasting and Social Change*. Additionally, as shown in Table 1.1, the rate of relevant article publication is increasing significantly, with more than 100 new articles published every year since 2018 and similarly more than 500 since 2020.

The above articles are classified into one or more of the following literature streams: financial theory, financial applications, statistical models, behavioral finance, crowdfunding, other financial topics, economics topics, financial misconduct,

⁷Note that considering English language documents yields an initial sample of 2,453 articles, reduced to 2,270 after removing less relevant articles. The following Scopus search query replicates these results: TITLE-ABS-KEY (("bitcoin" OR "Bitcoin" OR "ethereum" OR "cryptocurrency" OR "cryptocurrencies" OR "cryptoasset" OR "crypto asset" OR "digital currency" OR "crypto currency")) AND (LIMIT-TO (DOCTYPE,"ar")) AND (LIMIT-TO (SUBJAREA,"ECON") OR LIMIT-TO (SUBJAREA,"BUSI")) AND (LIMIT-TO (PUBYEAR,2019) OR LIMIT-TO (PUBYEAR,2018) OR LIMIT-TO (PUBYEAR,2017) OR LIMIT-TO (PUBYEAR,2016) OR LIMIT-TO (PUBYEAR,2015) OR LIMIT-TO (PUBYEAR,2014) OR LIMIT-TO (PUBYEAR,2013) OR LIMIT-TO (PUBYEAR,2020) OR LIMIT-TO (PUBYEAR,2021) OR LIMIT-TO (PUBYEAR,2022))

Table 1.1: Overview of the crypto asset and blockchain finance academic literature

Year	Count	%	Cumulative %
2013	1	0.04%	0.04%
2014	3	0.1%	0.2%
2015	24	1.1%	1%
2016	30	1.3%	3%
2017	44	2%	4%
2018	146	6%	11%
2019	374	16%	27%
2020	514	23%	50%
2021	629	28%	78%
2022	505	22%	100%

Distribution of articles by year in absolute, relative and cumulative relative count, based on a total of 2,270 articles published between 2013 – August 2022 on crypto asset and blockchain finance topics in finance, economics, econometrics, business, management and accounting journals.

regulatory topics, technology topics, and also other topics. The literature streams are structured based on a total of 168 topics, as follows.

Financial theory includes topics on asset pricing such as factor models to explain and/or predict crypto asset expected returns, price efficiency, market microstructure and price discovery, crypto asset option pricing and risk premia estimation.⁸ Financial applications include topics such as portfolio diversification with crypto assets, hedging, crypto asset arbitrage, and risk management.⁹

Statistical models include topics such as the examination of crypto asset returns stylized facts and distribution analysis, and volatility modelling and forecasting.¹⁰ Behavioral finance includes topics such as speculative investment behaviour, investor

⁸Other relevant topics in the stream of financial theory include: crypto asset valuation via prospect theory, network value estimation based on Metcalfe’s law and studies on causality.

⁹Other relevant topics in the stream of financial applications include: safe-haven and store-of-value properties of crypto assets, correlation with gold, market risk and forecasting of value-at-risk and expected shortfall, credit risk, volatility contagion in crypto markets, event studies, crypto asset fund management, crypto market crashes, crypto asset market liquidity, trading strategy price performance, and crypto asset interest rates and convenience yield.

¹⁰Other relevant topics in the stream of statistical models include: extreme value theory and heavy-tailed distribution applications, tail dependence models, extreme returns clustering, intra-day price distribution characteristics, returns jumps and mean-reversion, returns second moment characteristics, stochastic volatility models, volatility connectedness and spillover models, cointegration and co-movement models, econophysics applications, endogeneity studies, price bubble detection, and time series analysis models such as impulse response functions, the non-linear autoregressive distributed lag model, quantile autoregression models and the Taylor effect.

sentiment and attention and investor herding and herd effects.¹¹ Crowdfunding includes primarily topics on token offerings, initial coin and exchange offerings, and blockchain project financing by venture capital firms. Financial misconduct includes topics on market manipulation and criminal activities.¹² Other financial topics include stablecoins, non-fungible crypto asset tokens (NFTs), and decentralized finance (DeFi).¹³

The economics stream includes topics such as macroeconomic and monetary policy and economic integration of crypto assets.¹⁴ Regulatory topics include the effect of regulatory approaches and policy uncertainty on crypto assets, insurance, tax and accounting. Technology topics include crypto asset transaction costs, confirmation time and double spending vulnerabilities, and transaction privacy and anonymity.¹⁵ Finally, the ‘other topics’ stream includes topics such as digital art, digital forensics, health and Covid-19 related studies, and energy and supply chain applications.¹⁶

Table 1.2 exhibits the distribution of published articles across the 11 literature streams discussed above. Financial applications, statistical models and technology topics are the most highly-populated streams, with more than 600 articles included in each case. Note that the same article may belong to multiple streams, which is the case for approximately 33% of the articles examined. The multitude of topics and literature

¹¹Other relevant topics in the stream of behavioral finance include: investor ambiguity aversion, investor behavioral intention, characteristics and financial literacy of bitcoin users, and social media and social networks monitoring, and impact of behavioral factors on crypto asset returns.

¹²Crypto asset-related criminal activities include: fraud and corruption, money laundering, darknet and shadow economy activity, ransomware attacks and terrorist and hate group financing via crypto assets

¹³The ‘other financial topics’ stream also includes: lending and microfinance, the effect of crypto assets on banking and financial intermediation, central bank digital currencies (CBDCs), crypto asset-based fan tokens, blockchain applications on Islamic and Shariah-compliant finance

¹⁴Other topics included in the economics stream include cash substitution by crypto assets, currency competition, central banking and inflation.

¹⁵Other technology topics are: blockchain and full node analysis, blockchain splits (forks), blockchain governance and data security, consensus algorithm performance and design, hash function performance, distributed ledger system efficiency, crypto asset mining network analysis, smart contract applications, transaction identification via user graph analysis, crypto asset adoption and electronic cash acceptance, and crypto asset implementations as a payment system and medium of exchange.

¹⁶The ‘other topics’ stream also includes augmented and virtual reality and internet of things applications, analysis of crypto asset usage by social collectives, bibliometric and topic modelling, case studies and narrative analysis, e-government and e-voting, education and job market studies, tourism and land registration applications.

Table 1.2: Literature streams in crypto asset and blockchain finance

Literature Stream	Count	%
Crowdfunding	119	5%
Economics topics	151	7%
Financial misconduct	154	7%
Other financial topics	242	11%
Behavioral finance	348	15%
Regulatory topics	380	17%
Other topics	487	21%
Financial theory	494	22%
Technology topics	633	28%
Statistical models	864	38%
Financial applications	939	41%

Number of articles in absolute and relative count classified into streams in the crypto asset and blockchain finance literature, based on a total of 2,270 articles published between 2013 – August 2022 in finance, economics, econometrics, business, management and accounting journals; note that articles may belong to multiple literature streams.

streams identified as part of this very broad survey indicates clearly the wide variety of research topics covered by relevant academic articles, even when confining the survey to articles published in finance, economics, econometrics, business, management and accounting journals. For instance, extending the parameters of this survey to further include conference papers and also subject areas such as computer science, engineering and mathematics yields upwards of 12,000 articles, indicative of a cutting-edge multidisciplinary research area with fewer than 15 years in existence.

Thesis Contributions and Structure

This doctoral thesis contributes to the crypto asset and blockchain empirical finance literature in three key areas: (i) market risk modelling, developing simple volatility models which exhibit equal forecasting ability in terms of crypto asset tail risk measure and volatility forecasts, when compared against complex models; (ii) market manipulation, by extending a methodology derived from securities fraud litigation studies to identify blockchain transactions with a manipulative effect on crypto asset prices; (iii) crowdfunding via token offerings, by identifying factors of fundraising

success using regression models, and exploring how these factors vary across time. Each of the above contributions is developed in a separate chapter, as follows.

Chapter 2 on market risk modelling provides extensive backtests of hourly and daily Value-at-Risk (VaR) and Expected Shortfall (ES) forecasts regarded as best practice in the industry and used for regulatory approval. We test several exponentially-weighted moving average (EWMA) volatility and covariance model specifications and also introduce an asymmetric AEWMA model. We compare these specifications against an even simpler equally-weighted moving average random walk benchmark model and also against more complex univariate and multivariate GARCH specifications.

The forecasting accuracy of these models is assessed using data of both daily and hourly frequency for bitcoin, ether, ripple and litecoin. We examine both the left- and right-tail VaR and ES forecasts, with right-tail risk measures rarely covered in the relevant literature. We backtest VaR and ES with the industry standard (Basel Committee, 1996) traffic light test, as well as with more commonly-used VaR and ES backtesting methods. We also extend the ES traffic light test to cover right-tail ES backtesting. The models' volatility and covariance forecasting performance is also examined in terms of the entire distribution with scoring rules – the continuous ranked probability score (CRPS) for univariate distributions and the energy and variogram scores for multivariate forecasts.

The empirical results of this Chapter demonstrate that simpler models in the EWMA class are just as accurate as GARCH models for VaR and ES forecasting, and similarly when using average scores generated from proper univariate and multivariate scoring rules.

Chapter 3 on market manipulation examines large blockchain transactions of the tether stablecoin originating from the tether treasury and assesses whether they produce positive abnormal returns for bitcoin prices. The methodology is adapted from single-firm event studies used in securities fraud litigation, using regression factor models. This methodology can be useful in determining materiality and

producing damages estimates in legal cases of crypto asset market manipulation, as damages reports in securities legal cases are often rejected unless an event study-based methodology is used.

The hypothesis of significant positive abnormal bitcoin returns is examined with both parametric and non-parametric tests for robustness. The linear regression model used controls for the effect of past events in the estimation period and an index model specification is used instead of the constant-mean returns approach employed in the relevant literature. We further introduce a flexible event window size approach to account for clustering in USDT outflow events from the tether treasury, and exclude token platform swap transactions from the sample, as they are not related to market movements.

Additionally, we provide a detailed description of the blockchain data retrieval process from the Bitcoin/Omni Layer blockchain, based on the experience of setting up and operating a blockchain node locally. This is an important contribution, as only a small number of articles retrieve blockchain transaction data directly from the corresponding blockchains, and none of these articles provide detailed explanations on how this is achieved.

Chapter 4 on the tokenomics of crowdfunding examines the fundraising success of token offerings for the 2017 – early 2022 period, constituting one of the most comprehensive studies in this topic. We proxy fundraising success with the amount of funding raised and also by the minimum funding target exceedance. Success factors are derived from venture, token and offering characteristics, as well as additional common factors such as the price level of ether and the launchpad platforms used.

The fundraising success determinants introduced in this Chapter include variables that are not used in the relevant literature, such as the combination of token offering supply and price to form the target market capitalization, and also the launchpad platforms used by token offerings in 2021 – early 2022; we examine variables which are used in the relevant literature – such as the holding of a token presale, but for which the findings produce a lack of consensus as to the effect on fundraising success.

We also provide alternate perspectives on commonly-used variables such as the token offering rating and tax-haven domicile.

The chapter’s empirical results confirm and extend the findings in the relevant literature for the earlier sample period 2017 – 2020. Notably, this study is – at the time of writing, the first to identify the renewed activity of token offerings in 2021 – early 2022, with new all-time highs in the number of token offerings completed per month observed in late 2021 and early 2022. This period of renewed activity is accompanied by the emergence of initial decentralized exchange offerings (IDOs) which have dominated the token offerings space, as well as a shift in the fundraising success determinants, with the choice of launchpad platform emerging as a significant success determinant in 2021 – early 2022.

Finally, Chapter 5 presents a summary of the results for each topic developed together with relevant conclusions, and provides thoughts on potential further research in crypto asset and blockchain finance.

CRYPTO ASSET MARKET RISK MODELLING

2.1	Introduction	21
2.2	Literature Review	23
2.2.1	Models	27
2.2.2	In-Sample Performance	31
2.2.3	Out-of-Sample Performance	33
2.2.4	Focused Discussion and Contributions	38
2.3	Methodology	42
2.3.1	Volatility and Covariance Models	43
2.3.2	Value-at-Risk and Expected Shortfall	46
2.3.3	Distribution Forecast Evaluation	52
2.4	Data	55
2.5	Empirical Results	59
2.5.1	Daily Forecast Evaluation	59
2.5.2	Hourly Forecast Evaluation	71
2.6	Summary and Conclusions	80

2.1 Introduction

The modelling and forecasting of volatility and quantile risk measures for crypto assets is a fairly well-researched topic; just in the past year there are more than 100 academic papers published in this area of finance and similarly there are almost 350 papers when examining the past 8 years.¹ Common modelling choices include several variants from the generalised autoregressive conditional heteroscedasticity (GARCH) family of models initially introduced by Bollerslev (1986) and also other models such as the generalized autoregressive score (GAS) models introduced by Creal et al. (2013) and mixture and regime switching specifications. A similar degree of variety and complexity exists in the distribution assumptions for crypto asset returns: while the normal distribution is used by some authors, the most common choices are heavy-tailed distributions such as the Student- t ; moreover, several papers employ even more complex heavy-tailed and skewed distributions, such as the generalized error distribution (GED), the Weibull, Beta, generalized hyperbolic, inverse Gaussian and Johnson's SU distribution.

However, this complexity in modelling choices for crypto asset volatility and returns in the academic literature is in stark contrast with current practice among retail investors in crypto asset markets; they sometimes do not even apply any risk management apart from stop-loss limit orders placed at arbitrary price levels for open positions.² Moreover, the online sources that do discuss, use or provide forecasts of volatility, Value-at-Risk (VaR) and/or Expected Shortfall (ES) mostly use very

¹A relevant Scopus search yields 342 papers published between 2015 - early 2022 in Economics, Econometrics, Finance, Business, Management or Accounting journals and 131 of these papers were published in 2021 or early 2022. These results are produced with the following Scopus search query: TITLE-ABS-KEY(("bitcoin" OR "Bitcoin" OR "ethereum" OR "ether" OR "Ethereum" OR "Ether" OR "cryptocurrency" OR "cryptocurrencies" OR "cryptoasset" OR "crypto asset" OR "crypto" OR "digital currency" OR "digital asset" OR "crypto currency") AND ("GARCH" OR "EWMA" OR "Value at Risk" OR "VaR" OR "Value-at-Risk" OR "ES" OR "Expected shortfall" OR "volatility" OR "covariance" OR "variance") AND ("modelling" OR "modeling" OR "model" OR "models" OR "forecasting" OR "forecast" OR "forecasts" OR "estimation" OR "estimate" OR "estimates")) AND (LIMIT-TO(DOCTYPE, "ar")) AND (LIMIT-TO(SUBJAREA, "ECON") OR LIMIT-TO(SUBJAREA, "BUSI")).

²Note that the above refers to relatively unsophisticated retail investors that maintain unhedged positions in crypto assets; it does not apply e.g. to investors or market makers that partially or completely hedge their positions with derivatives.

simple methodologies. For instance, [Cryptodatadownload](#), a crypto asset market data and analytics provider, produces daily 1% and 5% VaR and ES forecasts for several crypto assets using a historical methodology over a 2-year period, i.e. the VaR is forecasted as the corresponding quantile of the empirical returns' distribution and ES as the average of the returns' observations that are lower than the corresponding quantile. Also, a blog post by the crypto asset exchange [OKEx](#) presents a parametric VaR estimation for bitcoin, under the assumption that its 1-minute returns follow a normal distribution; the 1% and 5% VaR are then forecasted using the sample mean and standard deviation of 1-minute returns over the past 7 days.³ Similarly, the daily '[Bitcoin Volatility Index](#)' is also calculated using the sample standard deviation of returns over the past 30 and 60 days; and the bitcoin [Fear & Greed Index](#) and a Forbes article (Bovaird, 2021) reporting on bitcoin's volatility both appear to be estimating volatility with a similar equally-weighted moving average.

Moreover, the estimation of volatility models such as GARCH requires the availability of historical price data for a fairly long period.⁴ While some crypto assets such as bitcoin or ether have been trading for long enough, the continuous emergence of new coins and tokens that gain investor attention often means that for some newer crypto assets, there may simply not be enough data available to produce robust parameter estimates for GARCH models. For instance, at the time of writing, the list of top crypto assets by market cap reported by [Cryptocompare](#) includes Avalanche and Solana which have just over 1.5 years of available data and also Terra which has only been trading for approximately 6 months. For such assets, volatility models that can be 'jump-started' and produce forecasts without the need for a lengthy estimation period, such as the RiskMetricsTM exponentially-weighted moving average (EWMA) model (Longerstaey and Spencer, 1996), are ideal.

³Moreover, the OKEx blog mentions that 'VaR is useful for calculating the maximum expected loss on an investment', which is a highly inaccurate interpretation of VaR, since VaR provides an estimate of *minimum* loss over the risk horizon.

⁴The optimal length of the sample period to be used in GARCH model estimation is highly dependent on the characteristics of the data used. For crypto asset volatility modelling, GARCH model estimations using between 1-2 years of daily data is a common choice in the relevant literature, given the frequent need for out-of-sample testing as well.

In this chapter, we compare the out-of-sample performance of parametric volatility and covariance models for bitcoin, ether, ripple and litecoin. In the following, Section 2.2 presents an overview of the relevant literature on crypto asset volatility and covariance forecasting using parametric models, VaR and ES backtesting and also crypto asset returns distribution forecast evaluation; in Section 2.3 we analyse the models used and also the backtesting and forecast evaluation methodologies; Section 2.4 provides an overview of the daily and hourly historical data; Section 2.5 presents the empirical results of VaR and ES backtesting for both the left and right tail of the returns' distribution and univariate and multivariate distribution forecast evaluation; finally, Section 2.6 provides a summary of the results and conclusions.

2.2 Literature Review

In this section, we provide a summary of the academic literature on crypto asset market risk modelling, focusing on papers that model the volatility and/or covariance of crypto asset returns with parametric models and assess their in-sample and out-of-sample performance. For ease of reference, the main characteristics of the most relevant academic papers are summarized in Table 2.1. For each paper, we report the assets examined and the sample period, the volatility and covariance models used and also the distribution assumptions in each case and the in- and out-of-sample methods used e.g. to determine the goodness-of-fit of competing models.

Table 2.1: Literature on crypto asset volatility and covariance forecasting

Author	Assets	Sample period	Models	Distributions	In-sample	Out-of-sample
Bouoiyour and Selmi (2016)	BTC	2011-2016	GARCH, EGARCH	Normal	AIC	HQ
			APARCH, wCGARCH		BIC	
			CMT GARCH			

Author	Assets	Sample period	Models	Distributions	In-sample	Out-of-sample
Chu et al. (2017)	BTC, XRP	2014-2017	GARCH, IGARCH	Normal	AIC	
	LTC, XMR		GJR-GARCH, EGARCH	Student- <i>t</i>	cAIC	
	DASH		APARCH, CGARCH	GED, SU	corAIC	
	DOGE		TGARCH, AVGARCH	gen. hyperbolic	BIC	
	MAID		NGARCH, AGARCH	inv. Gaussian	HQ	
	ALL GARCH					
Katsiampa (2017)	BTC	2010-2016	GARCH, APARCH	Normal	AIC	
			CGARCH, ACGARCH		BIC	
					HQ	
Baur et al. (2018)	BTC	2010-2015	EGARCH	Normal Student- <i>t</i>	Parameters	
Bonello and Suda (2018)	BTC	2016-2018	GARCH	Normal	AIC	UC, CC
			MS-GARCH	Student- <i>t</i>	BIC	
					DIC	
					Parameters	
Ardia et al. (2019)	BTC	2011-2018	GARCH, GJR-GARCH	Normal	DIC	CC, DQ
			MS-GARCH	Student- <i>t</i>		
			MS-GJR-GARCH			
Caporale and Zekokh (2019)	BTC, ETH XRP, LTC	2010-2018	Single-regime/mix./MS	Normal		UC, CC, DQ
		2013-2018	GARCH, GJR-GARCH	Student- <i>t</i>		ER, ESR
		2015-2018	EGARCH, TGARCH	GED		MCS
Catania et al. (2019)	BTC, ETH XRP, LTC	2015-2017	EWMA	Normal		MSE
			TVP-VAR			Log score
						MCS
Guesmi et al. (2019)	BTC	2012-2018	GARCH, GJR-GARCH	Normal	AIC	
			EGARCH, FIGARCH		BIC	
			FIAPARCH, DCC			
			ADCC, cDCC, cADCC			
Sosa et al. (2019)	BTC	2010-2019	GARCH, EGARCH	Normal	LL	
			TGARCH, APARCH	GED	AIC	
			CGARCH, ACGARCH		HQ	
Tiwari et al. (2019)	BTC	2011-2018	GARCH, GJR-GARCH	Normal	MLR	
	LTC		stochastic vol.	Student- <i>t</i>	Parameters	

Author	Assets	Sample period	Models	Distributions	In-sample	Out-of-sample
Trucíos (2019)	BTC	2011-2017	GARCH, AVGARCH GAS, GARCH-MIDAS realised-GARCH robust-GARCH	Normal Student- <i>t</i> GED, SU gen. hyperbolic inv. Gaussian		UC, CC, DQ MSE, QLIKE, RLF MCS
Troster et al. (2019)	BTC	2010-2018	EGARCH, GJR-GARCH APARCH, TGARCH CGARCH, NGARCH HGARCH, GAS	Normal Student- <i>t</i> GED, SU	AIC BIC	UC, CC, DQ RMSE
Wang et al. (2019)	BTC	2013-2018	GARCH, EGARCH CGARCH, ARJI	Normal	LL	Regression test
Acereda et al. (2020)	BTC, ETH XRP, LTC	2010-2018 2013-2018 2015-2018	GARCH, CGARCH NGARCH, TGARCH	Asymmetric Student- <i>t</i>		Multi-level
Alexander and Dakos (2020)	BTC	2013-2019	MS-GARCH MS-GJR-GARCH MS-EGARCH	Normal Student- <i>t</i>	DIC IC Parameters	
Bazán-Palomino (2020)	BTC, LTC BCH, BTG BCD, BTCP	2013-2019 2017-2019 2018-2019	EWMA BEKK-GARCH DCC-GARCH	Normal	Parameters Residuals	
Fantazzini and Zimin (2020)	BTC, ETH XRP, LTC XLM, eq. w. portfolio	2016-2018	GARCH DCC-GARCH copulas	Normal Student- <i>t</i>		UC, CC ER, Multi-level MCS
Hattori (2020)	BTC	2016-2018	GARCH, IGARCH GJR-GARCH EGARCH, APARCH	Normal Student- <i>t</i>		MSE QLIKE
Köchling et al. (2020)	BTC	2015-2018	GARCH, IGARCH GJR-GARCH, EGARCH APARCH, CGARCH AVGARCH, TGARCH NGARCH, AGARCH	Normal Student- <i>t</i>		MSE, MIX, QLIKE MCS
Liu et al. (2020)	BTC ETH LTC	until 2019	score-driven EWMA	Normal Student- <i>t</i> Laplace gen. Pareto reflected Gamma		UC, CC, DQ MCS

Author	Assets	Sample period	Models	Distributions	In-sample	Out-of-sample
Nekhili and Sultan (2020)	BTC, XRP	2014-2019	EWMA	Normal Student- <i>t</i>	LL AIC	CC, QL ER
	LTC, DASH		TGARCH			
	XMR, XLM		SVCJ			
	BCN					
Segnon and Bekiros (2020)	BTC	2013-2018	GARCH, GJR-GARCH	Normal		RMSE, MAE
			EGARCH, APARCH			MCS
			FIGARCH			LR
			MS-GARCH			
Catania and Grassi (2021)	606	until 2019	GAS EGARCH	Student- <i>t</i>	BIC	DQ
	large-cap			gen. hyperbolic		ER
	coins			Beta-skew- <i>t</i>		MSE, QLIKE
						CRPS
Maciel (2021)	BTC, ETH	2013-2018	Single-regime/MS	Normal	DIC	CC, DQ, QL
	XRP, LTC	2014-2018	GARCH, EGARCH	Student- <i>t</i>		FZL joint
	XMR, DASH	2015-2018	TGARCH	GED		DM
Silahli et al. (2021)	BTC, XRP	2014-2019	Hist. VaR	Normal Weibull		UC, CC, DQ
	LTC, DASH,		EQMA			
	min. var.		EWMA			
	portfolio		GARCH			

Key characteristics of the relevant academic papers that assess the forecasting performance of crypto asset volatility and covariance models. The columns indicate the author of each paper, the crypto assets examined and the sample period, the models and distribution assumptions used, and the in- and out-of-sample analysis performed.

As shown in Table 2.1, the crypto assets examined in relevant academic papers are: bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), dogecoin (DOGE), dash, monero (XMR), maidsafecoin (MAID), stellar (XML), bytecoin (BCN), bitcoin cash (BCH), bitcoin gold (BTG), bitcoin diamond (BCD), bitcoin private (BTCP), and also an equally-weighted and a minimum variance portfolio.⁵ The majority of papers focus on bitcoin, ether, ripple and litecoin, which offer a historical period of at least 5 years, as ether began trading in mid-2015. The above assets are also consistently amongst the largest cap crypto assets, so they can be considered as

⁵Historical data for the above assets are obtained by the authors from the following sources: blockchain.com, Binance, Bitstamp, Bloomberg, [Brave New Coin](https://bravenewcoin.com), [Coindesk](https://coindesk.com), Coinmarketcap, Cryptocompare, Gemini and Kraken.

representative of the entire asset class. However, other authors examine a more expanded crypto asset universe, such as Catania and Grassi (2021) who include in their analysis a total of 606 crypto assets with at least 700 daily price observations until September 2019. The sample frequency is almost invariably daily and the sample period used in each paper often depends on the available historical data for the assets included; for instance, Katsiampa (2017) and Baur et al. (2018) only examine bitcoin, so their sample period begins in 2010, whereas Fantazzini and Zimin (2020) use a sample period between 2016-2018, as they include ether in their analysis.

Having provided an overview of the crypto assets examined and sample periods used by papers in the relevant literature, we now discuss in turn: the parametric volatility and covariance models used and the corresponding distribution assumptions, followed by the results from the in-sample and out-of-sample performance analysis; finally we focus specifically on a small number of highly relevant papers and provide the incremental contribution of this chapter.

2.2.1 Models

The vast majority of academic papers examining crypto asset volatility employ some variant of the generalised autoregressive conditional heteroscedasticity (GARCH) model introduced by Bollerslev (1986), with a normal or Student- t distribution assumption for the returns.⁶

For instance, Dyhrberg (2016) uses a symmetric and an exponential normal GARCH with explanatory variables on bitcoin returns to compare bitcoin with gold and the dollar. Bouri et al. (2017) examine the hedging and safe-haven properties of bitcoin and use a symmetric model with innovations that follow a generalised error distribution (GED). Al-Khazali et al. (2018) compare the impact of macroeconomic news on bitcoin and gold attempting several model specifications and find that the optimal model is the exponential GARCH with normally distributed error terms.

⁶In the following, all GARCH-type models should be assumed to be of first order such as a GARCH(1,1), unless otherwise stated.

Corbet et al. (2018) examine the applications of bitcoin futures and use a symmetric GARCH. Vidal-Tomás and Ibañez (2018) use a component GARCH to examine the efficiency of bitcoin traded prices. Al-Yahyaee et al. (2019) study the diversification effects of bitcoin and gold for crude oil and S&P 500 investments and use several GARCH models including a fractionally integrated (FI) EGARCH model. López-Cabarcos et al. (2020) analyse the effect of investor sentiment and S&P 500 and VIX returns on bitcoin’s volatility using GARCH and EGARCH models.

Regarding the particular stream of the literature that examines the forecasting performance of parametric crypto asset volatility models, as presented in Table 2.1, the standard modelling choices include the symmetric GARCH of Bollerslev (1986) and asymmetric models such as the GJR-GARCH of Glosten et al. (1993), the exponential GARCH (EGARCH) of Nelson (1991), the threshold GARCH (TGARCH) of Zakoian (1994), the asymmetric power ARCH (APARCH) of Ding et al. (1993) and, less often, the AGARCH of Engle and Ng (1993). These models are in some cases extended further with distribution mixture and Markov switching (MS) frameworks.⁷

Further to the above, certain authors use the component GARCH (CGARCH) of Engle and Lee (1999) and variants such as its asymmetric extension ACGARCH, the weighted component GARCH (wCGARCH) of Bauwens and Storti (2009) and the component with multiple threshold (CMT) GARCH of Bouoiyour and Selmi (2014). Still more complex volatility model choices include the H-GARCH and ALL-GARCH of Hentschel (1995), the non-linear NGARCH of Higgins and Bera (1992), the AVGARCH of Schwert (1990), the robust GARCH model of Trucíos et al. (2017), the realised GARCH model of Hansen et al. (2012), the GARCH-MIDAS (mixed data sampling) model of Engle et al. (2013), and also an autoregressive jump intensity (ARJI) model and a stochastic volatility model with co-jumps (SVCJ). Similarly, more complex distribution assumptions – beyond the normal and Student- t and

⁷In most cases, the models are estimated using maximum likelihood estimation (MLE), with Markov chain Monte Carlo (MCMC) used in some cases for estimating regime-switching models; an exception to this are Tiwari et al. (2019) who estimate their models using the cross-entropy method of Rubinstein (1997) for calculating marginal likelihood.

their skewed variants, include the generalized error distribution (GED), generalized hyperbolic, Weibull, Laplace, Beta-skew- t , generalized Pareto, reflected Gamma, inverse Gaussian and Johnson’s SU distributions.

After discussing the more complex end of the volatility modelling spectrum, we now turn to simpler models such as RiskMetrics-type models which, by comparison, are somewhat overlooked in the relevant literature. The RiskMetrics™ EWMA model of Longestaey and Spencer (1996) is quite popular in financial market applications due to its simplicity and ease of use. As a result, some academic papers focus on assessing the RiskMetrics™ model’s forecasting accuracy using traditional asset data. For instance, Pafka and Kondor (2001) examine the VaR forecasting ability of the RiskMetrics™ model using the 30 constituent stocks of the Dow Jones Industrial Average equity index and argue that it performs well at higher significance levels and for short-term risk horizons, but that its accuracy declines e.g. at the 1% level and also for multi-period forecasts. Similarly, McMillan and Kambouroudis (2009) examine 31 stock market indices and note that the RiskMetrics™ model produces more accurate forecasts compared with GARCH at higher significance levels.⁸

Moreover, and specifically in the crypto asset financial literature, there is some support for the use of integrated volatility models such as the EWMA: for instance, Chu et al. (2017) and Köchling et al. (2020) examine an IGARCH model and find that it provides the optimal in-sample fit for bitcoin and other crypto assets; similarly, Bouoiyour and Selmi (2016) and Baur et al. (2018) find that bitcoin’s variance process is integrated while using GARCH-type models. The forecasting performance of EWMA volatility models is examined by Catania et al. (2019), Bazán-Palomino

⁸McMillan and Kambouroudis (2009) compare the 5% and 1% VaR forecasts of the RiskMetrics™ and several GARCH models using the unconditional coverage (UC) and dynamic quantile (DQ) tests of Kupiec (1995) and Engle and Manganelli (2004). At the higher significance level of 5%, the accuracy of RiskMetrics™ VaR forecasts is generally on par with the GARCH models used. However, the 1% VaR forecasts from the RiskMetrics™ model fail both backtesting processes. The authors further examine volatility forecasting accuracy with respect to realized volatility and find that more complex models approximate realized volatility more accurately than RiskMetrics™ in most cases examined.

(2020), Nekhili and Sultan (2020) and Silahli et al. (2021).⁹ Moreover, Liu et al. (2020) consider several score-driven EWMA models based on the generalized autoregressive score (GAS) model framework of Creal et al. (2013), while Trucíos (2019), Troster et al. (2019) and Catania and Grassi (2021) also use GAS models.

Finally, it is worth noting that the forecasting performance of multivariate covariance models is rarely examined and only in-sample, partly because the commonly-used out-of-sample VaR and ES backtesting methodology is better suited to a univariate setup. Bouri et al. (2017) are the first to examine crypto assets in a multivariate context, using a dynamic conditional correlation (DCC) model to test the hedge and safe-haven properties of bitcoin. The majority of relevant academic papers focus on in-sample performance and use the DCC model of Engle (2002), while some also make use of the earlier BEKK model of Engle and Kroner (1995). For instance, Bazán-Palomino (2020) considers the relationship between bitcoin and similarly structured crypto assets using the multivariate EWMA, BEKK-GARCH and DCC-GARCH, while Guesmi et al. (2019) use the DCC model to examine bitcoin as well as a number of traditional financial assets. Regarding the multivariate EWMA model, Matkovskyy et al. (2020) also use it to examine the interdependence between bitcoin, economic policy uncertainty and traditional financial assets, but none of the relevant papers assess its forecasting performance for crypto assets. Other covariance modelling choices reported in Table 2.1 include the asymmetric ADCC model of Cappiello et al. (2006), the modified cDCC and cADCC of Aielli (2013), and also multivariate extensions of the marginal densities using copula functions to model the correlation structure and time-varying parameter vector autoregression (TVP-VAR) models.

⁹Silahli et al. (2021) also examine an even simpler equally-weighted moving average (EQMA) model as a benchmark, while Guesmi et al. (2019) and Segnon and Bekiros (2020) use fractionally integrated models such as the FIGARCH and FIAPARCH.

2.2.2 In-Sample Performance

Following the discussion of volatility and covariance modelling choices for crypto assets, we now turn to the assessment of forecasting performance, starting with in-sample analysis. This is usually performed using the Akaike (AIC), Bayesian (BIC) and Hannan-Quinn (HQ) information criteria and also modified versions such as the consistent AIC (cAIC) and corrected AIC (corAIC); other in-sample performance criteria used in the relevant literature and shown in Table 2.1 include the direct comparison of the log likelihood (LL), the marginal likelihood ratio (MLR) and also, for models estimated via MCMC, the deviance information criterion (DIC) and the Bayesian predictive information criterion (IC). Beyond these, certain authors also examine the estimated model parameters and residuals to further illustrate the characteristics of the assets examined.

Numerous papers explore the best specification for univariate GARCH models on different types of financial data, see Engle et al. (2008) for a useful survey. Regarding the choice of volatility models providing the optimal in-sample fit to the crypto asset historical data examined, the results discussed below are often mixed and in contradiction with each other; this suggests that the best in-sample fit highly depends on both the assets and historical period examined in each case.

For instance, Bouoiyour and Selmi (2016) find that the APARCH and CMT GARCH provide the best in-sample fit for bitcoin returns, depending on the historical period examined. Chu et al. (2017) find mixed results on both the optimal model and distribution assumption, depending on the asset examined.¹⁰ Katsiampa (2017) also tests several parametric volatility models for the best in-sample fit on bitcoin returns and all criteria indicate that the ACGARCH model is optimal; this is consistent with Bouoiyour and Selmi (2016) whose in-sample analysis also indicates a model with a transitory and a permanent volatility component. The in-sample analysis of Baur et al. (2018) indicates the EGARCH model for bitcoin returns, and the

¹⁰The results of Chu et al. (2017) indicate that the normal distribution provides the best fit in all cases except ripple which ‘prefers’ a skewed normal distribution when using the TGARCH or the AVGARCH model. An IGARCH model is preferable for bitcoin, dash, litecoin, maidsafecoin and monero; GJR-GARCH is optimal for dogecoin; a symmetric GARCH is preferable for ripple.

authors note that using different asymmetric volatility models does not improve the in-sample fit. Tiwari et al. (2019) compare the fit of GARCH and stochastic volatility models for bitcoin and litecoin and find mixed results, concluding that crypto asset returns do not exhibit any asymmetric volatility response, which is at odds with the previous findings e.g. of Bouoiyour and Selmi (2016) and Katsiampa (2017). Sosa et al. (2019) find that an EGARCH model with GED innovations provides the best in-sample model fit for bitcoin, which is in contrast with the previous finding of Katsiampa (2017). Troster et al. (2019) agree with Sosa et al. (2019) that a GED assumption – instead of a normal, yields a significant improvement in the goodness-of-fit metrics, but they further conclude that the hyperbolic HGARCH model with GED innovations provides the best in-sample fit, which is again contrary to previous findings.

In the class of regime-switching volatility models, Ardia et al. (2019) find that a two-state Markov switching skewed Student- t GJR-GARCH provides a better in-sample fit for bitcoin as measured by the DIC, compared with both non-switching and three-state switching models; the authors propose that the two-state model provides a better trade-off between fitting quality and model complexity and further show for three-regime models that fitting gains are only observed for the normal distribution. Alexander and Dakos (2020) also explore the in-sample fit of two-state Markov switching GARCH models for bitcoin returns and show that the choice between GARCH, GJR-GARCH and EGARCH depends on the exact source of data used. More specifically, Alexander and Dakos (2020) examine the DIC and IC metrics for the daily log returns on BTC/USD price data from Coingecko, Coinmarketcap, Cryptocompare, Bitstamp and Kraken; Coingecko, Coinmarketcap and Bitstamp data ‘prefer’ a 2-state skewed Student- t standard GARCH model, while the 2-state skewed Student- t EGARCH provides a better in sample fit for Cryptocompare and Kraken data.

2.2.3 Out-of-Sample Performance

Out-of-sample forecasting exercises often focus on Value-at-Risk (VaR) and Expected Shortfall (ES) backtesting, usually concentrating on the left tail of the returns' distribution, i.e. to assess the risk of downward price movements on long crypto asset positions.¹¹ The more common backtesting methodologies for VaR forecasts are the unconditional coverage (UC) test of Kupiec (1995), the conditional coverage (CC) test of Christoffersen (1998) and the dynamic quantile (DQ) test of Engle and Manganelli (2004); for ES, common backtesting methods include the exceedance residual (ER) of McNeil and Frey (2000), the regression-based ESR test of Bayer and Dimitriadis (2020) and the multi-level backtest approximation via VaR of Kratz et al. (2018).¹² It is also worth noting that the industry standard traffic light backtesting framework of the Basel Committee (1996), e.g. as described by Costanzino and Curran (2018) is overlooked by all relevant papers. Other methods of analysis include the use of loss functions either in the model confidence set (MCS) process of Hansen et al. (2011) or also in hypothesis tests of equal forecasting performance such as the DM test of Diebold and Mariano (1995) and the regression test of Mincer and Zarnowitz (1969).¹³ Finally, the use of proper scoring rules to evaluate crypto asset returns density forecasts is much less common, with Catania and Grassi (2021) using the continuous ranked probability score (CRPS) and Catania et al. (2019) using the log score.

As noted previously for the in-sample analysis findings, the out-of-sample analysis results in the relevant literature discussed below are again conflicting and depend on the assets, sample periods and significance levels examined in each case. For instance, Ardia et al. (2019) examine bitcoin and compare the VaR forecasting

¹¹It is worth noting that the only relevant crypto asset paper that also assesses the performance of right-tail VaR and ES daily forecasts for bitcoin, i.e. the risk of upward price movements on short positions, is that of Stavroyiannis (2018), who examines the GJR-GARCH model.

¹²The test of Kratz et al. (2018) consists of a multinomial test of VaR exceptions at several levels below the significance level used for ES. See Section 2.3 for a detailed description of selected VaR and ES backtesting methodologies.

¹³Commonly-used loss functions include the MIX, quasi-like (QLIKE) and robust (RLF) loss functions of Patton (2011), the mean absolute error (MAE), mean squared error (MSE) and root mean squared error (RMSE), the quantile loss (QL) function of McAleer and Da Veiga (2008) and the FZL joint VaR/ES loss function of Fissler and Ziegel (2016).

accuracy of single-regime and Markov switching models using the CC and DQ tests, concluding that only regime-switching models produce accurate VaR forecasts at the 1% significance level; however, it is worth noting that 5% daily VaR forecasts produced using the relatively simpler single-regime skewed Student- t GJR-GARCH model also succeed the CC test – at the 5% significance level of rejecting the null hypothesis – and the DQ test. Maciel (2021) compares the prediction performance of Markov switching GARCH against single-regime GARCH models for several crypto assets and is in favour of more complex models – similar to Ardia et al. (2019), finding that overall, Markov switching models exhibit better results compared with the corresponding single-regime models; however, despite this finding, the results of Maciel (2021) are somewhat mixed.¹⁴ Caporale and Zekokh (2019) examine bitcoin, ether, ripple and litecoin and backtest both VaR using the UC, CC and DQ tests and ES using the ER and ESR backtests, also using the MCS process; their results are highly mixed due to the exhaustive use of mixture and regime switching model combinations, so that the distribution mixture symmetric GARCH is in some cases found to be preferable to asymmetric and Markov switching specifications.¹⁵

It further appears that even when more complex and even regime-switching volatility model specifications with heavy-tailed distribution assumptions produce accurate out-of-sample VaR and ES forecasts, relatively simpler models can also produce accurate results. For instance, Bonello and Suda (2018) also compare the VaR forecasts for bitcoin using the UC and CC tests with single-regime and Markov switching normal and Student- t GARCH models and find that both specifications can produce accurate VaR forecasts at a 5% significance level. Troster et al. (2019) use the UC, CC and DQ tests to backtest daily 1% VaR forecasts for bitcoin and

¹⁴For instance, Maciel (2021) finds no significant difference in forecasting performance between single- and two-regime models for half of the assets examined when forecasting performance is measured jointly for the 1-day-ahead 1% VaR and ES via the FZL loss function; similarly, the corresponding VaR-only test via the QL loss function indicates equal predictive ability between single- and two-regime models for all assets.

¹⁵The results of Caporale and Zekokh (2019) indicate that bitcoin and litecoin’s superior sets of models (SSM) is dominated by symmetric GARCH models with a slight preference for mixture models over Markov switching in the case of bitcoin; for ether, the only specification that produces accurate VaR and ES forecasts is a mixture of Student- t GJR-GARCH and skewed Student- t TGARCH, whereas the results for ripple are mixed.

find that a Student- t GARCH model succeeds in all of the above backtests and is therefore on par with more complex Student- t GARCH models attempted as well as with the attempted GAS model specification.¹⁶ Trucíos (2019) evaluates VaR forecasts for bitcoin between 2011 - 2017 and finds that out of six competing models, only a robust bootstrap VaR method produces accurate forecasts at the 1% significance level based on the UC, CC and DQ test. However, it should be noted that only a small selection of the proposed models are backtested based on the MCS process results.¹⁷ Moreover, in the preliminary results of a subsequent working paper, Trucíos and Taylor (2022) use a more recent sample period and show that bitcoin and ether VaR forecasts based on simpler volatility models such as the standard GARCH may be considered accurate based on the CC and DQ tests.¹⁸

There also exist several cases where backtesting results promote the use of even simpler models. For instance, Fantazzini and Zimin (2020) consider several crypto assets and backtest VaR and ES; interestingly, they find that conservative VaR forecasts are produced even when the normal distribution is used, and similarly for ES backtesting, the ER test does not reject the null hypothesis of accurate ES forecasts for almost any model.¹⁹ Acereda et al. (2020) find that bitcoin ES forecasts are accurate only when using a non-normal distribution with at least two parameters and more complex extensions of the standard GARCH model; for other crypto assets however, the more complex model specifications do not outperform the simpler ones, as long as heavy-tailed distributions are used instead of the standard normal. Silahli

¹⁶Beyond the standard GARCH model, Troster et al. (2019) also attempt the GJR-GARCH, EGARCH, CGARCH, NGARCH, and HGARCH models.

¹⁷Trucíos (2019) only selects to backtest the more complex model specifications that are included simultaneously in the MCS process superior sets of models based on four different realized volatility measures and three different loss functions; these models are the AVGARCH of Schwert (1990) with a symmetric and skewed GED assumption for innovations, the robust GARCH model of Trucíos et al. (2017) with a bootstrap VaR methodology, a Student- t GAS model and the realised GARCH model of Hansen et al. (2012).

¹⁸The working paper results of Trucíos and Taylor (2022) indicate that bitcoin 2.5% and 5% daily VaR forecasts produced e.g. with a standard GARCH model succeed in two out of three VaR backtesting methodologies and in all of the ES backtesting methodologies used. The only test that GARCH fails e.g. for bitcoin's 2.5% VaR forecasts is the VQ test of Gaglianone et al. (2011); for ether, the same is observed for 2.5% VaR forecasts; for ether's 5% VaR forecasts produced with the standard GARCH, all of the VaR and ES tests are successful.

¹⁹Fantazzini and Zimin (2020) attribute these findings to the small sample size (2016-2018) and the presence of excess kurtosis in the data.

et al. (2021) also find that several simple benchmark models succeed in the UC, CC and DQ VaR tests for several crypto assets.²⁰

Focusing on EWMA models, there is again conflicting evidence as to their forecasting accuracy. Silahli et al. (2021) find that the EWMA volatility model with a normal distribution assumption produces accurate VaR forecasts using UC, CC and DQ tests. Conversely, Liu et al. (2020) examine several crypto assets and find that a score-driven EWMA model specification similar to the standard RiskMetrics™ methodology (Longerstaey and Spencer, 1996) fails the backtesting processes and is often excluded from the MCS process superior set of models. Nekhili and Sultan (2020) compare the out-of-sample performance of stochastic volatility model specifications against a benchmark RiskMetrics™ EWMA model and find that the EWMA model produces accurate VaR forecasts at the 5% level, but not at 1%; for ES forecasts, the authors find that for almost all crypto assets examined, the EWMA model produces accurate ES forecasts according to the ER test.²¹

Further to the forecasting accuracy of EWMA models, in a multivariate setting, Silahli et al. (2021) find that a EWMA covariance model used to produce VaR forecasts for the return on a minimum variance portfolio consisting of bitcoin, litecoin, ripple and dash succeeds the UC, CC and DQ tests. Catania et al. (2019) examine bitcoin, ether, ripple and litecoin in a multivariate forecasting setting, testing several constant and time-varying parameter vector autoregression (VAR) models against a simpler VAR-EWMA benchmark. They find that the mean squared error (MSE) of forecasted volatilities against realized volatility and the MCS process using the MSE loss function show that none of the multivariate models can significantly outperform the VAR-EWMA benchmark at any forecast horizon; as discussed later,

²⁰For instance, Silahli et al. (2021) find that even at a 10% significance level, seven benchmark models succeed in the UC test, three succeed in the CC test and two succeed in the DQ.

²¹In the findings of Nekhili and Sultan (2020), the EWMA model's VaR forecasts are remarkably accurate for several crypto assets, given its simplicity compared with the other competing models. However, it should be noted the EWMA does not pass the CC test for the 5% VaR forecasts of litecoin and similarly for the 1% VaR of bitcoin, litecoin and stellar.

these findings are somewhat in contrast with the authors' results using the log score as a measure of density forecast accuracy.²²

Beyond VaR and ES backtesting, other out-of-sample methodologies also yield mixed results as to the optimal out-of-sample fit. Wang et al. (2019) examine bitcoin returns using a regression test against proxy measures of true volatility and find that EGARCH has the highest forecasting accuracy out of all GARCH-type models examined. Troster et al. (2019) compare the forecasts of various volatility models using RMSE measured against realized bitcoin returns and find that a CGARCH model with GED innovations provides the best out-of-sample forecasting accuracy. Köchling et al. (2020) also search for the best out-of-sample volatility model fit on bitcoin data using the MCS process and find that the IGARCH specifications are included in all MCS superior sets of models. Hattori (2020) finds that models with an asymmetric volatility response such as the EGARCH and APARCH exhibit higher predictive ability; surprisingly, Hattori (2020) further finds that models assuming a normal distribution perform better than models including a heavy-tailed distribution assumption. Segnon and Bekiros (2020) find that a Markov switching multi-fractal (MSM) model and also the FIGARCH model produce the most accurate volatility forecasts for bitcoin at both short- and long-term horizons but note that according to a likelihood ratio test, none of the models used can produce accurate density forecasts.

Finally, we examine proper scoring rules, a rarely-used out-sample analysis method even in the traditional finance literature. Regarding applications specific to crypto assets, the papers by Catania and Grassi (2021) and Catania et al. (2019) are the only relevant publications at the time of writing. Catania and Grassi (2021) use the continuous ranked probability score (CRPS) to assess volatility forecasts and compare the GAS model against an EGARCH, concluding that equal predictive ability as measured by the DM test of Diebold and Mariano (1995) is the most

²²For the EWMA model, Catania et al. (2019) assume a diagonal covariance matrix where variances are estimated with EWMA with $\lambda = 0.96$, i.e. a joint distribution for the four crypto assets with independent marginal densities.

common outcome. Moreover, Catania et al. (2019) produce multiperiod point and density forecasts for bitcoin, litecoin, ripple and ether returns with an expanding window and use the predictive log score as a measure of density forecast accuracy, which shows that most models outperform the EWMA benchmark, especially when including additional explanatory variables such as the VIX or equity indices; note that, as discussed previously, this finding is somewhat in conflict with the authors' findings using the mean squared error (MSE) of forecasted volatilities against realized volatility and the MCS process.

2.2.4 Focused Discussion and Contributions

Having examined the modelling approaches, the in-sample and out-of-sample performance methodologies and the findings of the relevant crypto asset literature, we now focus specifically on two highly relevant papers by Liu et al. (2020) and Catania and Grassi (2021), following which, we summarize the specific contributions of the current chapter. Liu et al. (2020) make extended use of EWMA-type models, including several extensions based on the GAS framework and the paper by Catania and Grassi (2021) is the only one using CRPS to assess the accuracy of density forecasts for crypto asset returns.

Liu et al. (2020) consider several score-driven generalizations of the RiskMetrics EWMA model to produce volatility and VaR forecasts for the daily log returns of bitcoin, ether and litecoin. The authors assume that crypto asset returns follow a normal, Student- t , Laplace, double generalized Pareto or reflected Gamma distribution, and volatility and other time-varying parameter dynamics are driven by the score of the forecasting distribution and estimated via MLE. The model variant closest to the RiskMetricsTM methodology of Longestaeey and Spencer (1996) is a EWMA variance model with a normal distribution assumption and decay parameter estimated via MLE in the GAS framework of Creal et al. (2013). VaR forecasts at 0.5%, 1%, 2.5% and 5% significance are calculated via the corresponding quantile function, depending on the distribution assumption. The authors evaluate the VaR forecasts

of each model using the UC, CC and DQ tests and also employ the model confidence set process. The VaR forecasts produced by the model closest to the RiskMetrics™ methodology fail most of the backtesting processes, especially at lower significance levels; regarding the MCS process, the authors find that the above RiskMetrics-like model ranks very low in the superior set of models (SSM) at lower VaR significance levels and, curiously, that it is excluded from the SSM for 5% VaR – contrary to all other relevant findings that simpler models produce reasonably accurate VaR forecasts at higher significance levels.

Catania and Grassi (2021) backtest VaR and ES forecasts for a total of 606 crypto assets with at least 700 daily price observations, i.e. approximately 2 years of historical data, until September 2019. The authors use score-driven (GAS) volatility model specifications that incorporate several stylized features such as leverage effects, long memory of the volatility process and time-varying higher order moments; it is further assumed that crypto asset returns follow a generalized hyperbolic skewed Student- t distribution as detailed in Aas and Haff (2006). As a benchmark model, the authors use the Beta-Skew- t -EGARCH of Harvey and Sucarrat (2014), under the rationale that score-driven models produce more accurate volatility estimates than GARCH in the presence of extreme observations. The model parameters are updated on a daily basis using an expanding window and 1-day-, 1-week- and 2-weeks-ahead VaR and ES forecasts are produced at the 5% and 1% significance levels. Value-at-Risk forecasts are backtested using the DQ test, while the ER test with 1,000 bootstrap replications is used for ES forecasts. The authors also use the continuous ranked probability score (CRPS) to assess the density forecasts of crypto asset returns.²³ They find that score-driven specifications produce accurate 5% and 1% ES and 5% VaR forecasts more often than the Beta-Skew- t -EGARCH benchmark, but GAS models and the EGARCH benchmark are on par when backtesting 1% VaR.²⁴

²³Catania and Grassi (2021) also test the models on the accuracy of their VaR, ES and also their variance predictions against the squared returns volatility proxy, using loss functions and the DM test of Diebold and Mariano (1995), and also report the average VaR absolute deviations and the resulting average daily capital requirements as per McAleer and Da Veiga (2008).

²⁴Catania and Grassi (2021) find that score-driven specifications produce accurate 5% VaR forecasts for approximately 70% of crypto assets examined, whereas the Beta-Skew- t -EGARCH

Regarding density forecast evaluation via CRPS and the test of equal forecasting performance, the authors find that outperformance of certain score-driven model specifications against the benchmark occurs more often than underperformance. However, even for these successful specifications, equal predictive ability is the most common outcome. For instance, when examining the uniformly weighted CRPS of the 1-day-ahead density forecast across all crypto assets, equal predictive ability occurs in 83% of crypto assets examined, including the large-cap bitcoin, ether, ripple and litecoin.

While both Liu et al. (2020) and Catania and Grassi (2021) examine several volatility model specifications, the range of models examined is somewhat limited in both cases. Liu et al. (2020) focus specifically on EWMA-type models and do not test other more complex models such as GARCH specifications, nor simpler model specifications that require no calibration such as an equally-weighted moving average or a EWMA with an ad hoc value chosen for the decay parameter. Therefore, their results are not conclusive with respect to the overall suitability of EWMA-type models in forecasting crypto asset volatility compared with other more complex or simpler models. By comparison, Catania and Grassi (2021) focus on highly sophisticated GAS model specifications with a similarly sophisticated heavy-tailed distribution assumption and test these against an already complex benchmark Beta-skew- t -EGARCH model. Their findings e.g. on density forecast accuracy often indicate that the benchmark is on par with the GAS specifications and in the absence of testing on simpler benchmark models, it is not always clear whether the additional modelling complexity ‘pays off’. It is important to note that, as discussed previously, the above finding also extends to VaR and ES forecasting, i.e. the VaR and ES forecasting performance of highly complex GARCH and GAS model specifications can be on par with relatively simpler models such as the standard GARCH. For

benchmark is accurate for approximately 60% of assets; however for 1% VaR, GAS models are considered accurate for approximately 55% of assets, whereas the EGARCH benchmark is still accurate for 60% of assets. Similarly, 5% ES forecasts are considered accurate for more than 90% of crypto assets examined and the Beta-Skew- t -EGARCH benchmark is accurate for 70% of assets, with slightly worse performance for 1% ES.

instance, this is shown in the results of Bonello and Suda (2018), Troster et al. (2019), Acereda et al. (2020), Silahli et al. (2021) and also in the working paper results of Trucíos and Taylor (2022).

Finally, regarding the contributions of this Chapter to the crypto asset financial literature, the main gaps in the area of crypto asset volatility and covariance forecasting with parametric models are located in the non-existent use of hourly-frequency data in volatility forecasting, the limited use of simpler and ad hoc models such as the EWMA, the very limited testing performed on right-tail risk measures, the complete absence of the industry standard traffic lights Value-at-Risk and Expected Shortfall backtesting methodology, the scant literature on the comparison of out-of-sample forecasting accuracy for multivariate models and the almost non-existent use of scoring rules for density forecast evaluation.

While the complex end of the modelling spectrum is well-researched, there is perhaps a limited scope for the application of such models from the perspective of a practitioner, especially that of an unsophisticated one such as retail investors. The EWMA model with ad hoc parameter choices is ideally suited for such use cases, as it is easy to understand, validate, explain in a simple technical document, and does not require long periods of historical data for calibration. Even when EWMA-type models are used, there is conflicting evidence in the relevant literature as to the accuracy of volatility, VaR, ES and density forecasts produced by them.

In this chapter, we test several EWMA model specifications with ad hoc parameter choices and introduce an asymmetric AEWMA model similar to the AGARCH model of Engle and Ng (1993). We compare these specifications against an even simpler equally-weighted moving average random walk benchmark model and also against more complex GARCH specifications, both in a univariate and a multivariate setting; a heavy-tailed Student- t returns distribution assumption is used for all models except the benchmark which assumes a normal distribution. We examine these models using historical price time series of both daily and hourly frequency for bitcoin, ether, ripple and litecoin, between August 2015 - August 2021 at the daily frequency and

between January 2021 - July 2021 at the hourly frequency; daily data frequency is the most common choice in the literature, but hourly data frequency is also well worth examining given how volatile crypto assets are, to the extent that an hour in the crypto asset market may almost be equivalent to an entire day in traditional financial markets.

Regarding VaR and ES, we examine both the left- and right-tail measures; right-tail VaR and ES backtesting are seldom covered in the relevant literature; given the wide availability of margin trading on crypto asset exchanges, the risk of upward price movements on short crypto asset positions can be measured via right-tail VaR and ES. To backtest VaR and ES, we make use of the traffic light test which is an industry standard (Basel Committee, 1996), as well as the more commonly-used CC test of Christoffersen (1998) for VaR and the ER test of McNeil and Frey (2000) for ES; the left-tail ES traffic light test is based on the methodology of Costanzino and Curran (2018), which we further extend to cover right-tail ES backtesting. Beyond risk measure forecast backtesting, we also examine the models' forecasting performance in terms of the entire distribution with scoring rules, using the continuous ranked probability score (CRPS) for univariate distributions and the energy and variogram scores for multivariate joint density forecasts.

2.3 Methodology

This Section provides an overview of the econometric methodologies used to measure the forecasting performance of the competing models. The overall methodology consists of producing 1-period-ahead volatility and covariance forecasts on a rolling basis. These are then combined with parametric distribution assumptions to produce forecasts for each asset's returns distribution and also of the joint density. We produce and backtest 1-period-ahead left- and right-tail Value-at-Risk and Expected Shortfall forecasts for each asset, to assess the risk of both long and short positions

on each asset. Similarly, 1-period-ahead distribution forecasts are evaluated via univariate and multivariate proper scoring rules.

2.3.1 Volatility and Covariance Models

The following models are used to forecast each asset’s volatility: a ‘random walk’ n -period equally-weighted moving average volatility model under a normal distribution assumption; a RiskMetrics-type EWMA model following from the RiskMetrics™ technical document (Longerstaey and Spencer, 1996) and also AEWMA, an asymmetric extension to EWMA similar to the AGARCH model of Engle and Ng (1993), both with a Student- t distribution assumption; the standard GARCH(1,1) model of Bollerslev (1986) and the asymmetric EGARCH model of Nelson (1991), again assuming that returns follow a Student- t distribution. Joint density forecasts are produced via the covariance matrix forecasts for the random walk and EWMA models. Similarly, the GARCH and EGARCH models are combined with the dynamic conditional correlation (DCC) model of Engle (2002) and Tse and Tsui (2002) and also its asymmetric extension (ADCC) model of Cappiello et al. (2006).

Univariate Models

Let the return of each asset at time t be r_t . For simplicity, it is assumed that the returns of all assets follow a zero-mean process.²⁵ The random walk benchmark model further assumes that the price of each asset follows a random walk process and therefore the return is:

$$r_t = \sigma_t \varepsilon_t \text{ and } \varepsilon_t \sim \mathcal{N}(0, 1), \quad (2.1)$$

where the volatility σ_t is estimated as the sample standard deviation of the returns over the past n periods. Similarly, in the EWMA and GARCH models presented

²⁵There is some support for the zero-mean assumption in the relevant literature as discussed previously in Section 2.2; for instance, Köchling et al. (2020) find that GARCH model specifications with a zero-mean assumption for bitcoin returns are very often included in the model confidence set.

below, returns are assumed to follow a zero-mean location-scale transformed Student- t distribution with ν degrees of freedom:

$$r_t = \sqrt{\frac{\nu - 2}{\nu}} \sigma_t \varepsilon_t \text{ and } \varepsilon_t \sim t_\nu, \quad (2.2)$$

where t_ν denotes that standardized Student- t distribution with ν degrees of freedom and σ_t is the standard deviation of r_t .

The variance under the standard EWMA model with decay parameter λ is calculated as:

$$\sigma_t^2 = (1 - \lambda)r_{t-1}^2 + \lambda\sigma_{t-1}^2. \quad (2.3)$$

Based on the AGARCH model of Engle and Ng (1993), we introduce the asymmetric EWMA model with a decay parameter λ and an asymmetric volatility response parameter η . Under the AEWMA(λ, η) model, the variance is calculated as:

$$\sigma_t^2 = (1 - \lambda)(r_{t-1} - \eta)^2 + \lambda\sigma_{t-1}^2. \quad (2.4)$$

In the standard (symmetric) GARCH(1,1) model, the conditional variance is given by:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2. \quad (2.5)$$

Similarly, in the Student- t EGARCH model, we have:

$$\begin{aligned} \ln(\sigma_t^2) &= \omega + g(\varepsilon_{t-1}) + \beta \ln(\sigma_{t-1}^2) \\ g(\varepsilon_t) &= \theta\varepsilon_t + \gamma(|\varepsilon_t| - \mathbb{E}[|\varepsilon_t|]). \end{aligned} \quad (2.6)$$

Regarding volatility forecasts, the random walk, EWMA and AEWMA models described in equations (2.1), (2.3) and (2.4) do not assume a volatility term structure, so their 1-period-ahead volatility forecast for time t is set equal to the corresponding conditional volatility estimate for time $t - 1$. For the GARCH and EGARCH

models the 1-period-ahead volatility forecast is obtained by updating the conditional volatility equations (2.5) and (2.6).

Multivariate Models

In a multivariate setting, let $\mathbf{r}_t$ be the $(m \times 1)$ vector of the m assets' returns at time t . The multivariate random walk benchmark model assumes that $\mathbf{r}_t$ follows a multivariate normal distribution:

$$\mathbf{r}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad (2.7)$$

where the covariance matrix $\boldsymbol{\Sigma}_t$ is estimated as the sample covariance matrix of assets' returns over the past n days. For the EWMA and GARCH models and following from their univariate counterparts, the vector of returns is assumed to follow a multivariate location-scale transformed Student- t distribution with ν degrees of freedom:

$$\mathbf{r}_t \sim t_\nu\left(\mathbf{0}, \frac{\nu-2}{\nu}\boldsymbol{\Sigma}_t\right), \quad (2.8)$$

where $\boldsymbol{\Sigma}_t$ is the covariance matrix of $\mathbf{r}_t$, so that $\frac{\nu-2}{\nu}\boldsymbol{\Sigma}_t$ is the distribution's scale matrix. The covariance matrix in the multivariate EWMA model with parameter λ is given by:

$$\boldsymbol{\Sigma}_t = (1 - \lambda)\mathbf{r}_{t-1}\mathbf{r}_{t-1}' + \lambda\boldsymbol{\Sigma}_{t-1}. \quad (2.9)$$

Similarly, the covariance matrix of the asymmetric EWMA with parameters λ and η is calculated as:

$$\boldsymbol{\Sigma}_t = (1 - \lambda)(\mathbf{r}_{t-1} - \eta\mathbb{1})(\mathbf{r}_{t-1} - \eta\mathbb{1})' + \lambda\boldsymbol{\Sigma}_{t-1}, \quad (2.10)$$

where $\mathbb{1}$ is an $(m \times 1)$ vector of ones.

In the DCC model and its asymmetric extension ADCC, the covariance matrix is modelled as:

$$\begin{aligned}\boldsymbol{\Sigma}_t &= \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t \\ \mathbf{R}_t &= \text{diag}(\mathbf{Q}_t)^{-1/2} \mathbf{Q}_t \text{diag}(\mathbf{Q}_t)^{-1/2},\end{aligned}\tag{2.11}$$

where $\mathbf{D}_t$ is the diagonal matrix of variances jointly estimated with univariate standard t GARCH or t EGARCH processes and $\mathbf{R}_t$ is the conditional correlation matrix, which is modelled indirectly via the $\mathbf{Q}_t$ matrix to ensure that the correlation matrix is positive semi-definite. In the DCC model, $\mathbf{Q}_t$ is given by:

$$\mathbf{Q}_t = (1 - a - b)\bar{\mathbf{Q}} + a\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}' + b\mathbf{Q}_{t-1}.\tag{2.12}$$

Similarly, in the ADCC model $\mathbf{Q}_t$ is calculated as:

$$\mathbf{Q}_t = (1 - a - b)\bar{\mathbf{Q}} - g\bar{\mathbf{Q}}^- + a\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}' + b\mathbf{Q}_{t-1} + g\boldsymbol{\varepsilon}_{t-1}^-\boldsymbol{\varepsilon}_{t-1}^{'-},\tag{2.13}$$

where $\boldsymbol{\varepsilon}_t$ is the vector of standardized errors, $\boldsymbol{\varepsilon}_t^-$ are the zero-threshold errors, $\bar{\mathbf{Q}}$ and $\bar{\mathbf{Q}}^-$ are the unconditional covariance matrices of $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\varepsilon}_t^-$.

The 1-period ahead covariance matrix forecasts are produced similar to the volatility forecasts as described previously. For the multivariate random walk, EWMA and AEWMA the 1-period-ahead covariance matrix forecast at time t is set equal to the estimate at time $t - 1$ and for the DCC and ADCC models it is obtained by updating the conditional covariance equation (2.11).

2.3.2 Value-at-Risk and Expected Shortfall

The forecasting accuracy of the volatility models presented in the previous Section is assessed by producing Value-at-Risk (VaR) and Expected Shortfall (ES) rolling forecasts and backtesting them against realized returns; VaR and ES forecasts are

produced for the left and right tail of the returns distribution, to assess the risk of both a long and short position on each asset. For each of the two quantile risk measures, we use an ad hoc methodology based on the standard traffic light test of the Basel Committee (1996) and also a standard backtesting method, i.e. the conditional coverage (CC) test of Christoffersen (1998) for VaR and the exceedance residual (ER) test of McNeil and Frey (2000) for ES.

VaR Backtesting

Value-at-Risk at a significance level α is defined as the α -quantile of the returns distribution. Therefore, based on the rolling 1-period-ahead forecast F_t of the returns' distribution for each asset at time t , the corresponding left- and right-tail VaR forecast at level α is:

$$\text{VaR}_t(\alpha) = \begin{cases} -F_t^{-1}(\alpha) & , \text{ for the left tail} \\ F_t^{-1}(1 - \alpha) & , \text{ for the right tail} \end{cases} \quad (2.14)$$

where F_t is the cumulative distribution function (CDF) of the 1-period-ahead distribution forecast for the return r_t of each asset, so that $F_t^{-1}(\alpha)$ is the $100\alpha\%$ quantile of the distribution.

The ad hoc backtesting method for VaR is based on the traffic light approach of the Basel Committee (1996), as described in Costanzino and Curran (2018), here extended to both left- and right-tail VaR. The exceedance indicator $X_t^{\text{VaR}}(\alpha)$ of each 1-period-ahead left- and right-tail $100\alpha\%$ -VaR forecast at times $t = 1, \dots, N$ is defined as:

$$X_t^{\text{VaR}}(\alpha) = \begin{cases} \mathbb{1}_{\{r_t \leq -\text{VaR}_t(\alpha)\}} & , \text{ for the left tail} \\ \mathbb{1}_{\{r_t \geq \text{VaR}_t(\alpha)\}} & , \text{ for the right tail} \end{cases} \quad (2.15)$$

where $\mathbb{1}_{\{\text{condition}\}}$ denotes an indicator function, i.e. equals 1 if the condition is satisfied and 0 otherwise. The cumulative number of VaR exceedances $X_N^{\text{VaR}}(\alpha)$ over the entire forecasting period $t = 1, \dots, N$ is then calculated as:

$$X_N^{\text{VaR}}(\alpha) = \sum_{t=1}^N X_t^{\text{VaR}}(\alpha). \quad (2.16)$$

Under the null hypothesis that the VaR model is specified correctly, the total number of VaR exceedances follows a binomial distribution with parameters N and α ;²⁶ we approximate the binomial with a normal distribution as:²⁷

$$X_N^{\text{VaR}}(\alpha) \sim \mathcal{N}(N\alpha, N\alpha(1-\alpha)). \quad (2.17)$$

Given the number of realized VaR exceedances x^{VaR} over the forecasting period, we also define the probability $\Phi(z)$ of obtaining x^{VaR} or fewer exceedances, where z is derived from the standard normal transformation of x^{VaR} and Φ is the CDF of the standard normal distribution.²⁸ The traffic light colour zones are then defined as: **Green** if $\Phi(z) < 0.95$; **Yellow** if $0.95 \leq \Phi(z) < 0.9999$; **Red** if $\Phi(z) \geq 0.9999$.

As described by the Basel Committee (1996), the three-zone approach is introduced to mitigate the statistical limitations of backtesting and balance the two error types: type I, i.e. the possibility that an accurate model is classified as inaccurate based on its backtesting results; type II, i.e. the possibility that an inaccurate model is not classified as such based on its backtesting results. In the green zone, the backtesting results are considered consistent with an accurate model and the probability of erroneously accepting an inaccurate model is low. In the red zone, the backtesting results are highly unlikely to have resulted from an accurate model, and the probability of erroneously rejecting an accurate model is low. In the

²⁶The null hypothesis that the VaR model is ‘specified correctly’ implies a joint hypothesis that the time series of VaR exceedance indicators $X_t^{\text{VaR}}(\alpha)$ is independent and identically distributed (i.i.d.) and that the proportion of realized VaR exceedances is equal to the VaR significance level α . Note that the above definition holds for both left- and right-tail VaR, as exceedances are defined respectively based on the $100\alpha\%$ left and right tail of the distribution.

²⁷The approximation of the binomial distribution with the normal is considered accurate based on the rule-of-thumb that both $N\alpha$ and $N(1-\alpha)$ should be greater than 5, which is the case for the analysis presented in Section 2.5, as we use $N = 1,704$ in the daily frequency analysis, $N = 1,465$ in the hourly frequency and $\alpha = 1\%$, 2.5% and 5% .

²⁸As the number of realized VaR exceedances x^{VaR} over the forecasting period is a realization of the random variable $X_N^{\text{VaR}}(\alpha)$ defined in equations 2.16 and 2.17, the probability of obtaining x^{VaR} or fewer VaR exceedances is given as: $P(X_N^{\text{VaR}}(\alpha) \leq x^{\text{VaR}}) = P\left(\frac{X_N^{\text{VaR}}(\alpha) - N\alpha}{\sqrt{N\alpha(1-\alpha)}} \leq \frac{x^{\text{VaR}} - N\alpha}{\sqrt{N\alpha(1-\alpha)}}\right) = \Phi(z)$, where $z \sim \mathcal{N}(0, 1)$.

yellow zone, backtesting results could be consistent with either accurate or inaccurate models, so additional information is required to determine whether the model used is specified correctly.

The VaR forecasts are further backtested using the conditional coverage (CC) test of Christoffersen (1998), for which the likelihood ratio test statistic LR_{cc} is:

$$LR_{cc} = \frac{\alpha^{n_1} (1 - \alpha)^{n_0}}{\left(\frac{n_{01}}{n_{00}+n_{01}}\right)^{n_{01}} \left(1 - \frac{n_{01}}{n_{00}+n_{01}}\right)^{n_{00}} \left(\frac{n_{11}}{n_{10}+n_{11}}\right)^{n_{11}} \left(1 - \frac{n_{11}}{n_{10}+n_{11}}\right)^{n_{10}}}, \quad (2.18)$$

where: α is the significance level used in the VaR model; n_1 is the number of realized VaR exceedances; $n_0 = N - n_1$ is the number of realized returns that do not exceed the VaR forecast; n_{00} is the number of non-exceedances preceded by a non-exceedance; n_{01} is the number of exceedances preceded by a non-exceedance; n_{10} is the number of non-exceedances preceded by an exceedance; n_{11} is the number of exceedances preceded by an exceedance.²⁹ The asymptotic distribution of $-2 \ln LR_{cc}$ is chi-squared with 2 degrees of freedom and the null hypothesis of the CC test is that $\left(\frac{n_{01}}{n_{00}+n_{01}}\right) = \left(\frac{n_{11}}{n_{10}+n_{11}}\right) = \alpha$, suggesting that the VaR forecasting model is specified correctly.

ES Backtesting

Expected Shortfall (ES) was introduced (Artzner et al., 1999; Acerbi and Tasche, 2002) to address the limitation of VaR in that it cannot capture tail risk beyond the specified quantile of the returns distribution (Basel Committee, 2012). Expected Shortfall is defined as the expected loss given that the corresponding VaR forecast is exceeded.³⁰ The 1-period-ahead ES forecast at level α for time t is calculated as the average VaR past the threshold α :

²⁹As in the case of the traffic light backtest, the conditional coverage test definitions hold for both left- and right-tail VaR, as exceedances are defined respectively based on the $100\alpha\%$ left and right tail of the distribution.

³⁰This definition of Expected Shortfall applies to both left- and right-tail VaR given the definition of left- and right-tail VaR exceedances in equation (2.15).

$$\text{ES}_t(\alpha) = \frac{1}{\alpha} \int_0^\alpha \text{VaR}_t(p) dp. \quad (2.19)$$

The ad hoc traffic light backtesting method for Expected Shortfall is introduced by Costanzino and Curran (2018), as a generalization of the VaR traffic light backtest of the Basel Committee (1996). Extending the discrete random variable representing VaR exceedances, Costanzino and Curran (2018) introduce the ES generalized exceedance indicator $X_t^{\text{ES}}(\alpha) \in [0, 1]$ for the left tail of the distribution; by applying the same methodology to right-tail VaR, we can define $X_t^{\text{ES}}(\alpha)$ as:

$$X_t^{\text{ES}}(\alpha) = \begin{cases} \left(1 - \frac{F_t(r_t)}{\alpha}\right) \mathbb{1}_{\{r_t \leq -\text{VaR}_t(\alpha)\}} & , \text{ for the left tail} \\ \left(1 - \frac{1-F_t(r_t)}{\alpha}\right) \mathbb{1}_{\{r_t \geq \text{VaR}_t(\alpha)\}} & , \text{ for the right tail} \end{cases} \quad (2.20)$$

where F_t is the CDF of the 1-period-ahead distribution forecast at time t for the return r_t of each asset.³¹ Note that, compared with $X_t^{\text{VaR}}(\alpha)$ defined in equation (2.15), $X_t^{\text{ES}}(\alpha)$ includes an additional term: $\left(1 - \frac{F_t(r_t)}{\alpha}\right)$ for the left tail and $\left(1 - \frac{1-F_t(r_t)}{\alpha}\right)$ for the right tail; this term determines the severity of each VaR exceedance, based on the realized loss beyond the VaR level.³² The cumulative ES generalized exceedance is then calculated as:

$$X_N^{\text{ES}}(\alpha) = \sum_{t=1}^N X_t^{\text{ES}}(\alpha). \quad (2.21)$$

Under the null hypothesis that the ES model is specified correctly, the distribution of $X_N^{\text{ES}}(\alpha)$ is provided by Costanzino and Curran (2018) based on the binomial

³¹The definition of the ES generalized exceedance indicator in equation (2.20) is derived by applying the definition of ES in equation (2.19) to the left- and right-tail VaR exceedance indicator $X_t^{\text{VaR}}(\alpha)$ defined in equation (2.15), i.e. $X_t^{\text{ES}}(\alpha) = \frac{1}{\alpha} \int_0^\alpha X_t^{\text{VaR}}(p) dp$.

³²Due to the severity term included in equation (2.20), realized returns that exceed the VaR forecast but not the ES forecast receive a relatively low weight and $X_t^{\text{ES}}(\alpha)$ is dominated by more severe exceedances, i.e. realized returns that most likely exceed both the VaR and ES forecasts. For instance, a left-tail VaR exceedance of high severity occurs for a very large negative realized return, so that $(1 - \frac{F_t(r_t)}{\alpha}) \approx 1$; conversely, for a left-tail VaR exceedance with low severity, the realized return is close to the 100%-VaR, so that $F_t(r_t) \approx \alpha$ and therefore $(1 - \frac{F_t(r_t)}{\alpha}) \approx 0$.

and Irwin-Hall distributions;³³ the authors further note that the distribution tends asymptotically to a normal distribution for large forecasting periods, based on the derivation of Costanzino and Curran (2015).³⁴

$$X_N^{\text{ES}}(\alpha) \sim \mathcal{N}\left(\frac{1}{2}N\alpha, N\alpha\left(\frac{4-3\alpha}{12}\right)\right). \quad (2.22)$$

Given the total realized ES generalized exceedances over the forecasting period x^{ES} , the probability of obtaining x^{ES} or fewer ES generalized exceedances is $\Phi(z)$, where z is again derived from the standard normal transformation of x^{ES} .³⁵ The traffic light colour zones are therefore again defined as: **Green** if $\Phi(z) < 0.95$; **Yellow** if $0.95 \leq \Phi(z) < 0.9999$; **Red** if $\Phi(z) \geq 0.9999$.

The ES forecasts are further backtested using the (raw) exceedance residual (ER) test of McNeil and Frey (2000). The backtest is based on the ES-specified residuals that exceed the VaR:

$$\varepsilon_t = \begin{cases} (-r_t - \text{ES}_t(\alpha)) \mathbb{1}_{\{r_t \leq -\text{VaR}_t(\alpha)\}} & , \text{ for the left tail} \\ (r_t - \text{ES}_t(\alpha)) \mathbb{1}_{\{r_t \geq \text{VaR}_t(\alpha)\}} & , \text{ for the right tail} \end{cases} \quad (2.23)$$

The ER test statistic is then calculated as the sample mean of ε_t :

$$\hat{\mu} = \begin{cases} \frac{\sum_{t=1}^N \varepsilon_t}{\sum_{t=1}^N \mathbb{1}_{\{r_t \leq -\text{VaR}_t(\alpha)\}}} & , \text{ for the left tail} \\ \frac{\sum_{t=1}^N \varepsilon_t}{\sum_{t=1}^N \mathbb{1}_{\{r_t \geq \text{VaR}_t(\alpha)\}}} & , \text{ for the right tail} \end{cases} \quad (2.24)$$

³³As noted previously for VaR, the null hypothesis that the ES model is ‘specified correctly’ implies a joint hypothesis that the time series of ES generalized exceedance indicators $X_t^{\text{ES}}(\alpha)$ is i.i.d. and that for all $p \in [0, \alpha]$, the probability of VaR exceedances is $P(r_t \leq -\text{VaR}_t(p)) = p$ for the left tail and $P(r_t \geq \text{VaR}_t(p)) = p$ for the right tail.

³⁴Note that the derivation described in Costanzino and Curran (2015) yields the asymptotic distribution of equation (2.22), whereas the distribution’s variance in equation (17) of Costanzino and Curran (2018) is mistakenly omitting a factor of N . Regarding the accuracy of the normal approximation, Clift et al. (2016) perform a simulation study and find that the approximation is accurate for a forecasting period of length $N = 250$; this confirms that the normal approximation is fit for use in our analysis, as we use forecasting periods that include over 1,000 observations in both the daily and hourly frequency analyses described in Section 2.5.

³⁵As noted previously for VaR, the total realized ES generalized exceedances x^{ES} over the forecasting period are a realization of the random variable $X_N^{\text{ES}}(\alpha)$ defined in equation (2.21) and the corresponding probability is again obtained via the standard normal transformation.

As the test statistic $\hat{\mu}$ does not have a standard distribution, we need to estimate its distribution using a bootstrap simulation.³⁶ The ER test's null hypothesis is that $\mathbb{E}[\varepsilon_t] = 0$; this is tested against a 1-sided alternative that $\mathbb{E}[\varepsilon_t] > 0$, suggesting that ES is systematically underestimated.

2.3.3 Distribution Forecast Evaluation

Scoring rules measure the accuracy of probabilistic forecasts and allow for comparisons between competing models. The models presented throughout this Section produce volatility and covariance 1-period-ahead forecasts which are combined with parametric distribution assumptions to forecast the marginal distributions and the joint density for the 1-period-ahead returns of the four crypto assets examined. We therefore use scoring rules to compare the accuracy of volatility and covariance forecasts between competing models.

A scoring rule is a function S such that:

$$S : \mathcal{F} \times \Omega \longrightarrow \mathbb{R} \cup \{\infty\} , \quad (2.25)$$

where $\mathcal{F}$ is a convex class of probability distributions on Ω . The scoring rule therefore assigns a value to a forecasted distribution $F \in \mathcal{F}$ and an observation $y \in \Omega$. Scoring rules are negatively oriented, such that a lower score indicates a better forecast. Moreover, for a scoring rule with the additional property of propriety, the expected score is optimized if the true distribution of the observation is issued as a forecast.

A scoring rule S is proper if and only if for all distributions F and G with densities f and g :

$$\int f(y)S(F, y)dy \leq \int f(y)S(G, y)dy . \quad (2.26)$$

We use the continuous ranked probability score (CRPS) for univariate distribution forecasts and its multivariate extension, the energy score, for joint density forecast evaluation. For multivariate forecasts, we also use the variogram score.

³⁶In the results presented in Section 2.5, the distribution of the ER test statistic $\hat{\mu}$ is simulated using 1,000 bootstrapped replications, similar to Catania and Grassi (2021).

Evaluation of Univariate Distribution Forecasts

The continuous ranked probability score (CRPS) (Matheson and Winkler, 1976 and Gneiting and Ranjan, 2011) generalizes the mean absolute error for the forecasted cumulative distribution function (CDF) F of a continuous distribution and a realized observation y :

$$\text{CRPS}(F, y) = \int_{-\infty}^{+\infty} \left(F(z) - \mathbb{1}_{\{y \leq z\}} \right)^2 u(z) dz \quad (2.27)$$

We choose a threshold weighting function $u(z) = 1$ to examine the entire distribution, but emphasizing specific parts of the distribution is also possible since any non-negative function can be used. Note that according to Gneiting and Raftery (2007), CRPS can also be expressed as:

$$\text{CRPS}(F, y) = \mathbb{E}_F(X - y) - \frac{1}{2} \mathbb{E}_F(X - X') , \quad (2.28)$$

where X and X' are independent random variables with sampling distribution F ; this representation leads to the energy score extension shown below.

Evaluation of Multivariate Distribution Forecasts

The energy score (**ES**) with unit index β generalizes the CRPS for multivariate distributions and is defined (Gneiting and Raftery, 2007) as:

$$\mathbf{ES}(F, \mathbf{y}) = \frac{1}{2} \mathbb{E}_F(\|\mathbf{X} - \mathbf{X}'\|) - \mathbb{E}_F(\|\mathbf{X} - \mathbf{y}\|) , \quad (2.29)$$

where $\|\cdot\|$ denotes the Euclidian norm on $\mathbb{R}^n$, $\mathbf{X}$ and $\mathbf{X}'$ are independent $(n \times 1)$ random vectors from a multivariate distribution with CDF forecast F and $\mathbf{y} = (y_1, \dots, y_n)$ is a realized observation. Moreover, if F is given via m discrete (n -dimensional) samples $\mathbf{X} = (X_1, \dots, X_n)$, then the energy score is calculated as:

$$\mathbf{ES}(F, \mathbf{y}) = \frac{1}{m} \sum_{i=1}^m \|\mathbf{X}_i - \mathbf{y}\| - \frac{1}{2m^2} \sum_{i=1}^m \sum_{j=1}^m \|\mathbf{X}_i - \mathbf{X}_j\| . \quad (2.30)$$

The variogram score of order p ($\mathbf{VS}^p$) is an alternative multivariate score (Scheuerer and Hamill, 2015) calculated as:

$$\mathbf{VS}^p(F, \mathbf{y}) = \sum_{i=1}^n \sum_{j=1}^n w_{i,j} \left(|y_i - y_j|^p - \frac{1}{m} \sum_{k=1}^m |X_{i,k} - X_{j,k}|^p \right)^2, \quad (2.31)$$

using $w_{i,j} = 1$ for all i, j , again to examine the entire multivariate distribution. We use the commonly suggested values $p = 0.5, 1$, and 2 . The variogram score with $p = 0.5$ is more sensitive to small deviations e.g. compared with $p = 2$ and conversely, the $p = 2$ variogram score tends to ‘punish’ large deviations more.

Comparison of Density Forecasts

Given the 1-period-ahead probability density function forecasts f_t, g_t and their corresponding univariate or multivariate scores $S(f_t)$ and $S(g_t)$ produced on a rolling basis over the out-of-sample period $t = 1, \dots, N$, we compare the forecasting performance of f and g directly using their average scores over the out-of-sample period. Alternatively, we use the hypothesis test of equal performance described by Gneiting and Ranjan (2011). If the average scores of f and g over the out-of-sample period are $\bar{S}_N^f$ and $\bar{S}_N^g$ respectively, then the test of equal performance is based on the test statistic

$$t_N = \sqrt{N} \frac{\bar{S}_N^f - \bar{S}_N^g}{\hat{\sigma}_N}, \quad (2.32)$$

where:

$$\hat{\sigma}_N^2 = \frac{1}{N} \sum_{t=1}^N \left(S(f_t) - S(g_t) \right)^2. \quad (2.33)$$

The test statistic t_N is asymptotically standard normal under the null hypothesis of vanishing expected score differentials; therefore in case of rejection, f is chosen if t_N is negative and g is chosen if t_N is positive.

2.4 Data

We obtain daily and hourly frequency historical data on four of the largest cap crypto assets as of 1 January 2021: bitcoin, ether, ripple and litecoin.³⁷ Historical price data are collected using the Cryptocompare API and are in the form of volume-weighted (VWAP) close prices, averaged across multiple USD-denominated exchange-traded prices for each crypto asset. Following from previous research (Alexander and Dakos, 2020), we choose to obtain the dollar-denominated VWAP price for each crypto asset from Cryptocompare, due to the advantages in its calculation methodology compared with other data sources.³⁸

For the daily frequency analysis, the sample period is between 20 August 2015 - 31 August 2021. The rolling estimation window length is chosen to be 500 days; the forecasting period therefore consists of 1,704 daily observations, between 1 January 2017 - 31 August 2021. For the hourly frequency analysis, the sample period is between 1 January 2021 00:00 UTC - 1 July 2021 00:00 UTC, with an estimation window length of 4 months, i.e. 2,882 hourly returns observations; the forecasting period therefore consists of 1,465 hourly observations, between 1 May 2021 00:00 UTC - 1 July 2021 00:00 UTC. Note that on one occasion for daily returns and on four occasions for hourly returns, the observations are considered as outliers and are removed via linear interpolation on the corresponding price.³⁹


Figure 2.1 shows the time-series of daily log returns for each crypto asset over the entire daily frequency sample period August 2015 - August 2021. Bitcoin appears to be considerably less volatile than the other assets with the exception of the ‘Black

³⁷We exclude stablecoins such as tether (USDT) from our analysis, as their prices are almost always very close to \$1 and therefore present limited interest from a volatility modelling perspective.

³⁸For instance, the Cryptocompare BTC/USD VWAP price is calculated using only the traded prices and corresponding volumes of the direct BTC/USD currency pair, contrary to other data sources that also employ a wide array of inferred prices via cross-rates and cross-rate volumes.

³⁹The daily data outliers are successive XRP returns of +103% on 2 April 2017 and -65% on 3 April 2017. The hourly data outliers are: an XRP return of -22% on 1 February 2021 13:00 UTC; XRP returns of 21% and -21% respectively on 22 February 2021 05:00 UTC and 23 February 2021 08:00 UTC; consecutive LTC returns of -35% and 22% on 19 May 2021 at 12:00 and 13:00 UTC. Outliers are detected based on the behaviour of the remaining asset returns and also on large return reversals indicative of deviations in some component of the volume-weighted average prices (VWAP) used.

Figure 2.1: Daily log returns of bitcoin, ether, ripple and litecoin



Daily log returns on bitcoin, ether, ripple and litecoin VWAP USD prices obtained from Crypto-compare. The sample period is 20 August 2015 - 31 August 2021.

Thursday' crypto market crash on 11 March 2020 and volatility clusters are often observed simultaneously across all four assets, as also confirmed the ARCH LM test statistic in Table 2.2.

Table 2.2 presents an overview of summary statistics for the daily log returns of bitcoin, ether, ripple and litecoin. All assets exhibit a relatively small mean, partly due to the use of log returns, so that a zero-mean assumption is justified.⁴⁰ The sample skewness always has a magnitude below 1, so the assumption of symmetrical

⁴⁰Given that the zero-mean assumption applies to all competing models used, the comparison of forecasting accuracy is not affected by this choice.

distributions for the returns is also justified. The high volatility (often above 100%) and positive excess kurtosis indicate that heavy-tailed distributions are to be preferred.

Table 2.2: Sample statistics of daily returns on bitcoin, ether, ripple and litecoin

	BTC	ETH	XRP	LTC
Mean (%)	0.242%	0.360%	0.228%	0.178%
St. Dev. (% p.a.)	77.1%	121.9%	143.9%	109.1%
Skewness	-0.83	-0.26	0.98	0.54
Ex. Kurtosis	12.07	6.81	11.41	13.04
Min.	-48%	-57%	-54%	-48%
Max.	23%	38%	62%	55%
LM	48.87	131.4	194.37	93.7


Summary statistics and 12th order ARCH LM test statistic of daily log returns on bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC) VWAP USD prices obtained from Cryptocompare. The sample period is 20 August 2015 - 31 August 2021. The mean, minimum and maximum are expressed in %, the daily standard deviation is annualized using a factor of $\sqrt{365}$.

Figure 2.2 shows the time-series of hourly log returns for each crypto asset over the entire sample period January - June 2021. All returns exhibit volatility clustering and extreme returns above 10% or below -10%, as also shown in the ARCH LM test statistic and the minimum and maximum returns in Table 2.3.

Table 2.3 shows the sample statistics of the hourly log returns for each crypto asset over the entire sample period January - June 2021. The mean and skewness are again relatively small, justifying the zero-mean and symmetrical distribution assumptions for hourly returns.⁴¹ As expected, hourly returns are more volatile compared with daily returns described in Table 2.2 and excess kurtosis is again positive, so a heavy-tailed distribution should be preferable.

⁴¹Interestingly, while the hourly returns exhibit a consistent negative skewness across all four assets, only the daily returns of bitcoin and ether follow this pattern.

Figure 2.2: Hourly log returns of bitcoin, ether, ripple and litecoin



Hourly log returns on bitcoin, ether, ripple and litecoin VWAP USD prices obtained from Crypto-compare. The sample period is 1 January 2021 - 1 July 2021.

Table 2.3: Sample statistics of hourly returns on bitcoin, ether, ripple and litecoin

	BTC	ETH	XRP	LTC
Mean (%)	0.004%	0.026%	0.028%	0.003%
St. Dev. (% p.a.)	104.5%	134.3%	183.6%	148.4%
Skewness	-0.23	-0.80	-0.28	-0.57
Ex. Kurtosis	9.84	8.78	8.10	5.67
Min. (%)	-11%	-17%	-18%	-13%
Max. (%)	12%	9%	13%	8%
LM	270.11	664.11	614.90	278.92

Summary statistics and 12th order ARCH LM test statistic of hourly log returns on bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC) VWAP USD prices obtained from Cryptocompare. The sample period is 1 January 2021 - 1 July 2021. The mean, minimum and maximum are expressed in % and the hourly standard deviation is annualized using a factor of $\sqrt{24 \times 365}$.

2.5 Empirical Results

In this Section we present and discuss the out-of-sample analysis for bitcoin, ether, ripple and litecoin daily and hourly log returns, comparing the results of the EWMA-type models against both the random walk benchmark and the more complex GARCH models. We begin with the daily frequency analysis and first present the backtesting of both left- and right-tail Value-at-Risk and Expected Shortfall forecasts in a univariate setting, followed by the evaluation both the univariate and multivariate density forecasts using scoring rules. The hourly frequency analysis is then presented in the same order.

2.5.1 Daily Forecast Evaluation

As discussed in sections 2.3 and 2.4, we produce 1-day-ahead volatility and covariance forecasts for the assets examined between 1 January 2017 - 31 August 2021, i.e. the forecasting period includes 1,704 daily observations. For the random walk benchmark model, forecasts are based on an equally-weighted 30-day moving average. EWMA and AEWMA volatilities and covariances are calculated using ad hoc values of 0.925 and 0.94 for the decay parameter λ , where 0.94 is the standard value proposed in the RiskMetricsTM framework; the AEWMA model further introduces an asymmetry

parameter η which is set to ad hoc values of 1%, 2% and 3% for left-tail (long position) VaR/ES forecasting and -1%, -3% and -5% for the right tail (short position).⁴² The univariate GARCH and multivariate DCC model parameters are estimated with maximum likelihood estimation (MLE) using a rolling estimation window of 500 daily observations and with model parameters updated on a daily basis.

Regarding the distribution assumptions, in the random walk benchmark model, crypto asset returns are assumed to follow a zero-mean normal distribution. In the EWMA and AEWMA models, a zero-mean location-scale transformed Student- t distribution is used with ad hoc $\nu = 6$ degrees of freedom, to produce a heavy-tailed distribution; similarly, a multivariate Student- t with $\nu = 6$ is assumed for the joint distribution of bitcoin, ether, ripple and litecoin returns. The GARCH and DCC models also assume respectively a univariate and multivariate Student- t distribution, where the degrees of freedom parameter is estimated jointly with the model parameters based on the 500-day rolling estimation window.⁴³ GARCH and DCC model estimations and forecasts and also some of the VaR and ES backtesting methods are implemented using the rugarch and rmgarch R packages of Ghalanos (2020) and Ghalanos (2019).⁴⁴

VaR and ES Backtesting

Value-at-Risk and Expected Shortfall 1-day-ahead forecasts are produced for bitcoin, ether and ripple at the 1%, 2.5% and 5% significance levels for both the left and right tail of the returns distribution, to assess the risk of downward price movements on long positions and also of upward price movements on short positions. VaR forecasts are backtested using the traffic light test as described by Costanzino and Curran

⁴²Additional testing of other AEWMA η parameter choices is shown in Table A1 of Appendix A for the 1% daily VaR forecasts and the corresponding traffic light test.

⁴³The daily returns exhibit some asymmetry in their sample statistics, but not always in the same direction, as shown in Table 2.2. The asymmetric model specifications can still capture this characteristic to some extent even with the symmetric distribution assumption that we use. Also, the choice of a symmetric distribution simplifies the process of modelling the joint density for the returns of bitcoin, ether, ripple and litecoin.

⁴⁴The multivariate EWMA model is implemented using the RiskPortfolios R package of Ardia et al. (2017), and the remaining EWMA and AEWMA specifications and also the traffic light backtesting methodology are implemented using custom-written R code.

(2018) and also the conditional coverage (CC) test of Christoffersen (1998). Similarly, we apply the generalized traffic light test of Costanzino and Curran (2018) for ES backtesting and also use the ER test of McNeil and Frey (2000).

VaR backtesting results at the 1%, 2.5% and 5% significance level are shown in Table 2.4 for the left tail (long position) and in Table 2.5 for the right tail (short position). For each asset we report the number of realized VaR exceedances x^{VaR} , the probability $\Phi(z)$ of obtaining x^{VaR} or fewer VaR exceedances under the normal distribution approximation described in equation (2.16) and the p-value of the CC test statistic calculated as per equation (2.18).

As shown in Table 2.4 (upper panel) for 1% left-tail VaR, the random walk benchmark produces significantly more VaR exceedances than the EWMA and GARCH models across all assets. However, it should be noted that nearly all of the models examined exceed the expected number of 17.04 exceedances for 1% VaR with the exception of two AEWMA specifications and the EGARCH.⁴⁵ The traffic light zones are defined as per the Basel Committee (1996) framework described by Costanzino and Curran (2018), based on the cumulative probability $\Phi(z)$ of obtaining x^{VaR} or fewer VaR exceedances: **Green** if $\Phi(z)$ is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. Based on these definitions, the VaR forecasts of several AEWMA models and also the of t EGARCH model are in the green zone for all assets, with several other specifications partly ‘in the green’. Moreover, the forecasts of several models are often in the yellow zone; as per the Basel Committee (1996) standard framework, backtesting results in the yellow zone could indicate an inaccurate VaR model, so further testing should be performed.

Further to the ad hoc traffic light test for 1% left-tail VaR forecasts, we now discuss the results of the conditional coverage (CC) test, where rejection of the null hypothesis suggests that the VaR model used is not accurate. At the 10% significance level for the rejection of the CC test’s null hypothesis of accurate VaR forecasts,

⁴⁵The expected number of total VaR exceedances for 1% VaR is calculated based on equation (2.17) using $N = 1,704$ and $\alpha = 1\%$.

several AEWMA specifications and also the EGARCH model appear to produce accurate 1-day-ahead 1% VaR forecasts simultaneously for all assets examined; it is also worth noting that e.g. the 1% VaR forecasts from the t AEWMA(0.925, 2%) model are in the yellow zone for bitcoin, but the CC test results suggest that they are in fact accurate.⁴⁶

By comparison, for the right-tail daily 1% VaR backtesting results shown in Table 2.5, the results are similar for bitcoin and ether forecasts, in that several AEWMA and both GARCH specifications are in the green zone. However, ripple and litecoin VaR forecasts present more difficulties, with several models in the red zone, possibly due to the positive sample skewness that ripple and litecoin exhibit as shown in Table 2.2. Even so, the forecasts of two AEWMA specifications are in the green zone for all four assets, when the 1% right-tail VaR forecasts produced by the more complex t EGARCH model are actually considered inaccurate for ripple, with a yellow zone and rejection of the CC test's null hypothesis at both 10% and 5% significance.

Similarly, Expected Shortfall backtesting results at 1%, 2.5% and 5% are shown in Table 2.6 for the left tail (long position) and in Table 2.7 for the right tail (short position). Here we report for each asset: the realized total ES generalized exceedances x^{ES} as defined in equations 2.20 and 2.21; the probability $\Phi(z)$ of obtaining x^{ES} of fewer ES generalized exceedances under the normal distribution approximation as described in equation (2.22); and the p-value of the exceedance residual (ER) test based on equation (2.24), where the test statistic's distribution is simulated using 1,000 bootstrapped replications.

As with 1% daily VaR examined previously, daily 1% left-tail ES backtesting results shown in Table 2.6 (upper panel) indicate that the random walk benchmark produces significantly more ES generalized exceedances than the EWMA and GARCH models. Here, the expected number of ES generalized exceedances at 1% significance

⁴⁶Note that as the CC test for VaR and the ER test for ES have a null hypothesis of accurate model specification, a lower significance level, e.g. 1% for the rejection of the null hypothesis means that we are less strict about accepting the VaR and ES models as accurate. Conversely, for a higher significance level, e.g. 10%, we are less strict about rejecting the null hypothesis of accuracy and therefore more strict about accepting VaR and ES models as accurate.

Table 2.4: Backtesting results for 1-day-ahead left-tail VaR

Daily 1% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	38	1.0000	0.000***	38	1.0000	0.000***	28	0.9962	0.011**	33	0.9999	0.001***
t EWMA(0.925)	27	0.9923	0.017**	28	0.9962	0.002***	22	0.8864	0.050*	24	0.9549	0.037**
t EWMA(0.94)	28	0.9962	0.002***	27	0.9923	0.002***	20	0.7644	0.383	23	0.9266	0.044**
t AEWMA(0.925, 1%)	30	0.9992	0.001***	28	0.9962	0.011**	20	0.7644	0.383	26	0.9854	0.023**
t AEWMA(0.925, 2%)	24	0.9549	0.180	22	0.8864	0.291	17	0.4961	0.842	21	0.8325	0.499
t AEWMA(0.925, 3%)	17	0.4961	0.373	19	0.6834	0.723	15	0.3097	0.770	17	0.4961	0.842
t AEWMA(0.94, 1%)	28	0.9962	0.002***	25	0.9737	0.132	19	0.6834	0.404	23	0.9266	0.235
t AEWMA(0.94, 2%)	23	0.9266	0.235	20	0.7644	0.383	16	0.4001	0.832	19	0.6834	0.723
t AEWMA(0.94, 3%)	16	0.4001	0.325	19	0.6834	0.723	15	0.3097	0.770	17	0.4961	0.842
t GARCH	26	0.9854	0.086*	17	0.4961	0.842	13	0.1627	0.534	19	0.6834	0.723
t EGARCH	16	0.4001	0.832	16	0.4001	0.832	13	0.1627	0.534	14	0.2296	0.665

Daily 2.5% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	58	0.9916	0.004***	63	0.9992	0.004***	54	0.9615	0.024**	58	0.9916	0.032**
t EWMA(0.925)	54	0.9615	0.070*	56	0.9812	0.005***	46	0.7011	0.107	54	0.9615	0.024**
t EWMA(0.94)	47	0.7526	0.112	55	0.9728	0.021**	43	0.5247	0.081*	50	0.8746	0.107
t AEWMA(0.925, 1%)	50	0.8746	0.107	54	0.9615	0.024**	44	0.5860	0.092*	51	0.9038	0.100*
t AEWMA(0.925, 2%)	42	0.4629	0.265	44	0.5860	0.092*	35	0.1191	0.454	39	0.2882	0.497
t AEWMA(0.925, 3%)	35	0.1191	0.212	36	0.1529	0.555	28	0.0117	0.042**	29	0.0174	0.066*
t AEWMA(0.94, 1%)	45	0.6452	0.329	54	0.9615	0.024**	41	0.4020	0.233	48	0.7990	0.113
t AEWMA(0.94, 2%)	41	0.4020	0.233	45	0.6452	0.101	32	0.0500	0.203	38	0.2377	0.131
t AEWMA(0.94, 3%)	32	0.0500	0.078*	33	0.0682	0.275	23	0.0012	0.003***	29	0.0174	0.066*
t GARCH	54	0.9615	0.154	39	0.2882	0.846	42	0.4629	0.995	47	0.7526	0.670
t EGARCH	37	0.1924	0.658	37	0.1924	0.346	33	0.0682	0.113	39	0.2882	0.846

Daily 5% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	87	0.5793	0.266	90	0.7032	0.027**	78	0.2118	0.170	85	0.4911	0.100*
t EWMA(0.925)	91	0.7404	0.074*	86	0.5354	0.113	78	0.2118	0.008***	83	0.4034	0.357
t EWMA(0.94)	87	0.5793	0.266	82	0.3610	0.063*	70	0.0456	0.007***	87	0.5793	0.484
t AEWMA(0.925, 1%)	88	0.6222	0.283	87	0.5793	0.052*	75	0.1284	0.034**	81	0.3203	0.279
t AEWMA(0.925, 2%)	77	0.1810	0.282	78	0.2118	0.026**	62	0.0050	0.002***	74	0.1066	0.152
t AEWMA(0.925, 3%)	61	0.0036	0.016**	65	0.0124	0.003***	49	0.0000	0.000***	62	0.0050	0.023**
t AEWMA(0.94, 1%)	88	0.6222	0.502	81	0.3203	0.018**	70	0.0456	0.021**	81	0.3203	0.130
t AEWMA(0.94, 2%)	76	0.1532	0.234	77	0.1810	0.057*	58	0.0012	0.001***	71	0.0572	0.029**
t AEWMA(0.94, 3%)	60	0.0025	0.006***	63	0.0068	0.010**	48	0.0000	0.000***	62	0.0050	0.014**
t GARCH	99	0.9375	0.002***	87	0.5793	0.125	85	0.4911	0.429	91	0.7404	0.710
t EGARCH	91	0.7404	0.030**	83	0.4034	0.612	73	0.0875	0.120	77	0.1810	0.139

Backtesting results for 1-day-ahead left-tail 1%, 2.5% and 5% VaR forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 January 2017 - 31 August 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the total number of realized VaR exceedances x^{VaR} ; the probability $\Phi(z)$ of obtaining the realized number of VaR exceedances; the p-value of the CC test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the CC test, the null hypothesis is that the VaR forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

Table 2.5: Backtesting results for 1-day-ahead right-tail VaR

Daily 1% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	37	1.0000	0.000***	50	1.0000	0.000***	57	1.0000	0.000***	44	1.0000	0.000***
t EWMA(0.925)	28	0.9962	0.040**	27	0.9923	0.017**	50	1.0000	0.000***	34	1.0000	0.000***
t EWMA(0.94)	28	0.9962	0.040**	28	0.9962	0.002***	53	1.0000	0.000***	33	0.9999	0.000***
t AEWMA(0.925,-1%)	22	0.8864	0.291	27	0.9923	0.002***	46	1.0000	0.000***	35	1.0000	0.000***
t AEWMA(0.925,-3%)	13	0.1627	0.534	15	0.3097	0.263	30	0.9992	0.000***	24	0.9549	0.037**
t AEWMA(0.925,-5%)	5	0.0017	0.003***	4	0.0007	0.001***	17	0.4961	0.000***	11	0.0707	0.002***
t AEWMA(0.94,-1%)	21	0.8325	0.343	28	0.9962	0.011**	47	1.0000	0.000***	33	0.9999	0.001***
t AEWMA(0.94,-3%)	12	0.1099	0.397	14	0.2296	0.197	30	0.9992	0.000***	22	0.8864	0.050*
t AEWMA(0.94,-5%)	4	0.0007	0.001***	4	0.0007	0.001***	18	0.5924	0.000***	10	0.0433	0.001***
t GARCH	23	0.9266	0.283	16	0.4001	0.325	34	1.0000	0.001***	26	0.9854	0.086*
t EGARCH	17	0.4961	0.842	13	0.1627	0.534	28	0.9962	0.032**	14	0.2296	0.665

Daily 2.5% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	59	0.9945	0.055*	74	1.0000	0.000***	85	1.0000	0.000***	66	0.9999	0.002***
t EWMA(0.925)	52	0.9277	0.091*	69	1.0000	0.000***	78	1.0000	0.000***	67	0.9999	0.002***
t EWMA(0.94)	49	0.8397	0.111	64	0.9996	0.000***	73	1.0000	0.000***	62	0.9987	0.010**
t AEWMA(0.925,-1%)	45	0.6452	0.329	59	0.9945	0.001***	75	1.0000	0.000***	64	0.9996	0.005***
t AEWMA(0.925,-3%)	24	0.0020	0.005***	37	0.1924	0.003***	58	0.9916	0.000***	45	0.6452	0.101
t AEWMA(0.925,-5%)	10	0.0000	0.000***	23	0.0012	0.000***	39	0.2882	0.000***	27	0.0077	0.007***
t AEWMA(0.94,-1%)	44	0.5860	0.735	56	0.9812	0.001***	69	1.0000	0.000***	61	0.9978	0.014**
t AEWMA(0.94,-3%)	23	0.0012	0.002***	37	0.1924	0.003***	58	0.9916	0.000***	41	0.4020	0.633
t AEWMA(0.94,-5%)	10	0.0000	0.000***	23	0.0012	0.000***	39	0.2882	0.000***	27	0.0077	0.007***
t GARCH	55	0.9728	0.148	49	0.8397	0.301	63	0.9992	0.007***	63	0.9992	0.011**
t EGARCH	40	0.3433	0.918	44	0.5860	0.315	57	0.9873	0.078*	54	0.9615	0.230

Daily 5% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	105	0.9861	0.011**	113	0.9990	0.000***	103	0.9761	0.000***	94	0.8360	0.455
t EWMA(0.925)	95	0.8620	0.166	118	0.9999	0.000***	106	0.9896	0.000***	97	0.9052	0.251
t EWMA(0.94)	89	0.6636	0.296	113	0.9990	0.000***	101	0.9605	0.000***	92	0.7751	0.499
t AEWMA(0.925,-1%)	82	0.3610	0.560	108	0.9944	0.001***	104	0.9817	0.000***	91	0.7404	0.511
t AEWMA(0.925,-3%)	42	0.0000	0.000***	63	0.0068	0.001***	87	0.5793	0.000***	71	0.0572	0.141
t AEWMA(0.925,-5%)	21	0.0000	0.000***	39	0.0000	0.000***	59	0.0018	0.000***	49	0.0000	0.000***
t AEWMA(0.94,-1%)	81	0.3203	0.505	100	0.9500	0.016**	100	0.9500	0.000***	89	0.6636	0.749
t AEWMA(0.94,-3%)	45	0.0000	0.000***	61	0.0036	0.002***	81	0.3203	0.000***	71	0.0572	0.141
t AEWMA(0.94,-5%)	20	0.0000	0.000***	41	0.0000	0.000***	57	0.0009	0.000***	47	0.0000	0.000***
t GARCH	121	1.0000	0.001***	109	0.9959	0.020**	109	0.9959	0.000***	115	0.9995	0.002***
t EGARCH	100	0.9500	0.276	96	0.8850	0.392	102	0.9691	0.001***	100	0.9500	0.276

Backtesting results for 1-day-ahead right-tail 1%, 2.5% and 5% VaR forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 January 2017 - 31 August 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the total number of realized VaR exceedances x^{VaR} ; the probability $\Phi(z)$ of obtaining the realized number of VaR exceedances; the p-value of the CC test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the CC test, the null hypothesis is that the VaR forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

is 8.52 and, contrary to VaR, none of the models produce fewer ES generalized exceedances than the expected value simultaneously for all four assets; however, several AEWMA models and both GARCH models produce fewer than 8.52 ES exceedances for left-tail 1% ES e.g. for ripple and litecoin.⁴⁷ The traffic light zones for ES are defined similar to VaR, based on $\Phi(z)$. According to the traffic light ES test, forecasts based on two AEWMA specifications and also the t EGARCH model are in the green zone for all assets examined, and the symmetric t GARCH performs almost equally well with the exception of bitcoin. In the exceedance residual (ER) test for ES, rejection of the null hypothesis suggests that the ES model used is not accurate. The 1-day-ahead 1% ES forecasts produced several AEWMA models and also the GARCH and EGARCH are considered accurate even at the 10% significance level of rejecting the ER test's null hypothesis.

By comparison, for the right-tail daily 1% ES backtesting results shown in Table 2.7, the results are similar for bitcoin and ether forecasts, i.e. multiple AEWMA and GARCH specifications are in the green zone. However, as observed in the case of 1% right-tail VaR, ripple and litecoin ES forecasts again present difficulties, with several models in the red zone. In spite of this, the forecasts of two AEWMA specifications are yet again in the green zone for all four assets, and succeed the ER test at 10% or 5% significance.

Further to 1% VaR and ES, the daily right- and left-tail VaR and ES backtesting results at 2.5% significance shown in tables 2.4, 2.5, 2.6 and 2.7 (middle panel) and at 5% (lower panel), indicate that even more EWMA-type model specifications can produce accurate forecasts at higher VaR and ES significance levels. For instance, the left-tail 2.5% daily VaR forecasts of the t AEWMA(0.94, 2%) model are in green and succeed the CC VaR backtest for all assets at the 10% significance level of null hypothesis rejection. Similarly, in the 5% VaR and ES forecasts, all models attempted, even the random walk benchmark are either in the green or yellow zone and very often fail to reject the null hypothesis of the CC and ER tests, indicating

⁴⁷The expected value of the total ES generalized exceedances for 1% ES is calculated based on equation (2.22) using $N = 1,704$ and $\alpha = 1\%$.

Table 2.6: Backtesting results for 1-day-ahead left-tail ES

Daily 1% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	29.0	1.0000	0.00***	28.2	1.0000	0.00***	22.1	1.0000	0.00***	25.0	1.0000	0.00***
t EWMA(0.925)	17.8	1.0000	0.08*	17.3	0.9999	0.11	14.8	0.9961	0.07*	15.1	0.9972	0.12
t EWMA(0.94)	17.6	0.9999	0.13	15.8	0.9990	0.13	12.9	0.9670	0.14	13.6	0.9833	0.15
t AEWMA(0.925, 1%)	18.3	1.0000	0.12	15.4	0.9982	0.19	13.5	0.9826	0.09*	14.2	0.9915	0.20
t AEWMA(0.925, 2%)	14.4	0.9935	0.23	13.2	0.9759	0.15	11.0	0.8534	0.16	11.6	0.9021	0.26
t AEWMA(0.925, 3%)	8.9	0.5623	0.36	11.5	0.8977	0.25	7.7	0.3664	0.27	9.0	0.5816	0.31
t AEWMA(0.94, 1%)	18.2	1.0000	0.09*	14.5	0.9939	0.20	11.9	0.9233	0.16	13.1	0.9733	0.19
t AEWMA(0.94, 2%)	14.2	0.9917	0.23	13.1	0.9726	0.14	10.0	0.7278	0.21	11.2	0.8720	0.18
t AEWMA(0.94, 3%)	8.3	0.4695	0.36	11.4	0.8867	0.27	6.8	0.2386	0.33	8.8	0.5496	0.28
t GARCH	12.6	0.9557	0.67	9.2	0.6134	0.55	7.5	0.3358	0.42	8.9	0.5690	0.69
t EGARCH	8.7	0.5369	0.61	8.3	0.4560	0.62	6.4	0.1866	0.65	7.5	0.3354	0.55

Daily 2.5% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	39.6	1.0000	0.00***	41.3	1.0000	0.00***	33.5	0.9995	0.00***	36.8	1.0000	0.00***
t EWMA(0.925)	30.8	0.9947	0.08*	32.8	0.9990	0.08*	25.5	0.8689	0.14	30.3	0.9923	0.10
t EWMA(0.94)	30.0	0.9898	0.04**	31.2	0.9959	0.12	23.1	0.6855	0.20	27.6	0.9555	0.13
t AEWMA(0.925, 1%)	30.9	0.9949	0.06*	30.8	0.9944	0.12	23.6	0.7346	0.18	28.1	0.9650	0.13
t AEWMA(0.925, 2%)	25.8	0.8881	0.10	26.2	0.9059	0.08*	19.1	0.2781	0.20	22.3	0.6069	0.11
t AEWMA(0.925, 3%)	18.8	0.2557	0.35	21.5	0.5171	0.13	15.2	0.0517	0.23	17.9	0.1826	0.10
t AEWMA(0.94, 1%)	30.0	0.9898	0.03**	29.6	0.9867	0.15	21.6	0.5336	0.21	26.1	0.9021	0.13
t AEWMA(0.94, 2%)	25.1	0.8448	0.12	25.2	0.8512	0.12	17.6	0.1603	0.23	21.1	0.4757	0.13
t AEWMA(0.94, 3%)	18.5	0.2280	0.28	20.5	0.4154	0.09*	14.2	0.0278	0.18	17.2	0.1373	0.14
t GARCH	27.8	0.9596	0.66	22.2	0.5930	0.36	20.2	0.3806	0.67	24.1	0.7695	0.62
t EGARCH	20.2	0.3844	0.66	19.9	0.3566	0.56	15.7	0.0654	0.72	18.7	0.2392	0.75

Daily 5% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	56.8	0.9967	0.00***	58.3	0.9987	0.00***	50.0	0.9215	0.00***	53.7	0.9831	0.00***
t EWMA(0.925)	52.7	0.9729	0.05*	50.7	0.9399	0.02**	44.1	0.6110	0.06*	49.9	0.9186	0.03**
t EWMA(0.94)	50.3	0.9304	0.07*	49.5	0.9074	0.02**	40.7	0.3580	0.07*	48.1	0.8527	0.10*
t AEWMA(0.925, 1%)	51.3	0.9512	0.07*	49.9	0.9185	0.04**	40.4	0.3393	0.11	47.7	0.8369	0.05**
t AEWMA(0.925, 2%)	43.0	0.5325	0.11	44.2	0.6169	0.06*	33.5	0.0414	0.14	40.4	0.3339	0.12
t AEWMA(0.925, 3%)	32.7	0.0298	0.27	36.0	0.1036	0.10*	26.5	0.0011	0.17	31.5	0.0167	0.22
t AEWMA(0.94, 1%)	49.8	0.9171	0.09*	48.7	0.8801	0.02**	37.7	0.1723	0.14	46.0	0.7406	0.07*
t AEWMA(0.94, 2%)	42.2	0.4698	0.12	43.5	0.5701	0.07*	31.5	0.0170	0.18	38.8	0.2337	0.14
t AEWMA(0.94, 3%)	32.0	0.0215	0.29	35.3	0.0822	0.09*	25.3	0.0005	0.26	30.4	0.0097	0.28
t GARCH	54.4	0.9881	0.46	42.9	0.5241	0.62	41.6	0.4268	0.68	47.3	0.8153	0.61
t EGARCH	42.0	0.4514	0.94	38.7	0.2299	0.74	34.7	0.0644	0.81	39.3	0.2644	0.78

Backtesting results for 1-day-ahead left-tail 1%, 2.5% and 5% ES forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 January 2017 - 31 August 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the realized total ES generalized exceedances x^{ES} ; the probability $\Phi(z)$ of obtaining the realized total ES generalized exceedances; the p-value of the ER test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the ER test, the null hypothesis is that the ES forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

Table 2.7: Backtesting results for 1-day-ahead right-tail ES

Daily 1% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	29.5	1.0000	0.00***	32.3	1.0000	0.00***	47.2	1.0000	0.00***	32.6	1.0000	0.00***
t EWMA(0.925)	17.4	0.9999	0.09*	17.5	0.9999	0.14	33.1	1.0000	0.01**	21.7	1.0000	0.07*
t EWMA(0.94)	16.1	0.9993	0.18	16.9	0.9998	0.25	31.5	1.0000	0.04**	20.4	1.0000	0.10
t AEWMA(0.925,-1%)	15.1	0.9971	0.06*	16.0	0.9992	0.29	31.9	1.0000	0.01***	19.8	1.0000	0.15
t AEWMA(0.925,-3%)	6.5	0.1979	0.73	7.0	0.2622	0.76	19.1	1.0000	0.03**	12.4	0.9474	0.35
t AEWMA(0.925,-5%)	1.7	0.0019	1.00	2.4	0.0050	0.90	11.6	0.8993	0.03**	5.9	0.1360	0.46
t AEWMA(0.94,-1%)	14.0	0.9899	0.08*	15.4	0.9982	0.43	30.4	1.0000	0.03**	18.5	1.0000	0.17
t AEWMA(0.94,-3%)	6.0	0.1441	0.69	6.4	0.1872	0.80	18.9	1.0000	0.02**	11.6	0.9014	0.31
t AEWMA(0.94,-5%)	1.6	0.0017	1.00	2.3	0.0044	0.93	11.5	0.8952	0.06*	6.0	0.1463	0.36
t GARCH	11.1	0.8568	0.90	7.2	0.2857	0.98	16.0	0.9992	0.81	12.5	0.9546	0.71
t EGARCH	6.8	0.2332	0.95	5.6	0.1061	0.99	13.0	0.9716	0.93	6.5	0.1960	0.71

Daily 2.5% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	40.0	1.0000	0.00***	51.0	1.0000	0.00***	62.0	1.0000	0.00***	47.0	1.0000	0.00***
t EWMA(0.925)	30.3	0.9921	0.06*	35.2	0.9999	0.49	53.5	1.0000	0.00***	41.4	1.0000	0.01**
t EWMA(0.94)	28.8	0.9780	0.08*	33.5	0.9995	0.45	51.1	1.0000	0.00***	38.9	1.0000	0.02**
t AEWMA(0.925,-1%)	25.9	0.8897	0.12	31.4	0.9966	0.44	51.0	1.0000	0.00***	38.2	1.0000	0.03**
t AEWMA(0.925,-3%)	13.0	0.0132	0.39	19.4	0.3086	0.72	35.2	0.9999	0.05**	24.9	0.8295	0.17
t AEWMA(0.925,-5%)	5.2	0.0000	0.79	8.4	0.0003	0.98	21.8	0.5488	0.23	13.6	0.0202	0.40
t AEWMA(0.94,-1%)	25.4	0.8648	0.14	30.5	0.9934	0.45	49.5	1.0000	0.00***	36.2	1.0000	0.05**
t AEWMA(0.94,-3%)	12.5	0.0090	0.41	19.0	0.2730	0.73	34.9	0.9999	0.04**	23.8	0.7461	0.16
t AEWMA(0.94,-5%)	5.0	0.0000	0.80	8.0	0.0002	0.99	21.5	0.5190	0.22	13.3	0.0158	0.42
t GARCH	27.0	0.9368	0.88	21.1	0.4813	1.00	36.8	1.0000	0.41	32.3	0.9983	0.66
t EGARCH	19.5	0.3195	0.97	17.5	0.1571	1.00	31.9	0.9977	0.67	23.2	0.6909	0.96

Daily 5% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	61.7	0.9999	0.00***	70.2	1.0000	0.00***	79.1	1.0000	0.00***	63.5	1.0000	0.00***
t EWMA(0.925)	53.6	0.9820	0.05*	63.9	1.0000	0.25	73.0	1.0000	0.00***	60.3	0.9997	0.00***
t EWMA(0.94)	51.3	0.9511	0.07*	60.9	0.9998	0.24	70.5	1.0000	0.00***	58.9	0.9991	0.00***
t AEWMA(0.925,-1%)	44.9	0.6702	0.14	56.7	0.9964	0.32	70.2	1.0000	0.00***	57.9	0.9983	0.00***
t AEWMA(0.925,-3%)	23.5	0.0001	0.30	34.4	0.0591	0.42	53.1	0.9774	0.00***	42.5	0.4902	0.04**
t AEWMA(0.925,-5%)	10.7	0.0000	0.65	20.6	0.0000	0.83	35.9	0.0986	0.03**	26.2	0.0009	0.26
t AEWMA(0.94,-1%)	43.9	0.5990	0.17	54.4	0.9883	0.26	68.1	1.0000	0.00***	56.5	0.9961	0.00***
t AEWMA(0.94,-3%)	23.3	0.0001	0.45	34.1	0.0525	0.38	51.3	0.9519	0.00***	41.5	0.4150	0.05*
t AEWMA(0.94,-5%)	10.4	0.0000	0.64	20.3	0.0000	0.91	35.5	0.0867	0.02**	25.4	0.0005	0.23
t GARCH	58.5	0.9988	0.96	49.1	0.8940	1.00	60.2	0.9996	0.30	61.1	0.9998	0.52
t EGARCH	45.2	0.6923	1.00	44.8	0.6596	1.00	55.8	0.9941	0.62	49.3	0.8993	0.91

Backtesting results for 1-day-ahead right-tail 1%, 2.5% and 5% ES forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 January 2017 - 31 August 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the realized total ES generalized exceedances x^{ES} ; the probability $\Phi(z)$ of obtaining the realized total ES generalized exceedances; the p-value of the ER test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the ER test, the null hypothesis is that the ES forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

that the forecasts are accurate. In fact, when forecasting 5% 1-day-ahead left-tail VaR, even the random walk benchmark, is in the green zone for all four assets and succeeds in the CC test at 10% significance for half of the assets, suggesting that producing 5% 1-day-ahead left-tail VaR and ES forecasts sometimes does not even require a heavy-tailed distribution assumption, so that the simple moving average volatility model is almost on par with the more complex GARCH and EGARCH models and with EWMA-type models.

The 1-day-ahead VaR and ES backtesting results shown in this Section are consistent with most of the findings in the relevant literature as discussed in Section 2.2, such as Liu et al. (2020) who focus on VaR backtesting and find that VaR forecasts produced by RiskMetrics-type models are increasingly accurate as the VaR significance level increases. Moreover, Catania and Grassi (2021) examine GAS model specifications against an EGARCH benchmark for 606 crypto assets and find that the additional modelling complexity introduced by the GAS framework ‘pays off’ for 5% and 1% ES and 5% VaR with increased accuracy, but less so for 1% VaR. In that respect, the results presented in this Section for 1-day-ahead VaR and ES forecast backtesting are somewhat in agreement with Catania and Grassi (2021) in that introducing additional modelling complexity may sometimes ‘pay off’ in increased forecasting accuracy, especially at lower significance levels; however, we often find that AEWMA specifications are on par with EGARCH in terms of VaR and ES forecasting accuracy even at the 1% significance level.

Distribution Forecast Evaluation

Having examined quantile risk measures, we now present the results on univariate and multivariate scores to measure the accuracy of the competing models at forecasting the entire distribution of crypto asset 1-day-ahead log returns. The continuous ranked probability score (CRPS) is used to assess univariate density forecasts and joint distribution forecasts are evaluated with the energy score and also with the variogram scores of orders $p = 0.5, 1$ and 2 . Given the parametric distribution assumptions in

the models used, i.e. normal for the random walk model and Student- t for the EWMA and GARCH specifications, the 1-day-ahead volatility and covariance forecasts fully define the 1-day-ahead distribution of log returns for each asset and also their joint distribution, allowing for the scores' calculation. All univariate and multivariate scores are calculated using the `scoringRules` R package of Jordan et al. (2019).

The CRPS is calculated via a closed form formula for the 1-day-ahead log returns density forecast produced by each model, based on the corresponding realized return. For comparison purposes, the CRPS is averaged across the 4.5-year forecasting period. Table 2.8 reports the outright average CRPS for the random walk benchmark model and the average scores of all other models are expressed as a percentage of benchmark model's average score. Due to the negative orientation of scoring rules, relative scores below 100% suggest possible outperformance against the benchmark, and vice-versa. Beyond an outright comparison of the average scores, we also use the scores to perform pair-wise comparisons of forecasting accuracy and calculate the t_N test statistic of Gneiting and Ranjan (2011) for the hypothesis test of equal forecasting performance as per equation (2.32) in Section 2.3.

As shown in Table 2.8, the random walk benchmark produces a higher average CRPS compared with most other models, although none of them achieve an average CRPS lower than 97% of the benchmark's average score. Amongst them, the $tAEWMA(0.94, -5\%)$ model produces the highest CRPS for its density forecasts across all four assets examined. The model with the lowest average score is always a GARCH specification: for bitcoin, ether and litecoin it is the $tEGARCH$ and for ripple it is the $tGARCH$. However, and more importantly, the t_N test statistic calculated between the highest and lowest average scores for each asset is always below 0.15, so the null hypothesis of equal forecasting performance is always accepted e.g. at the 5% and 10% significance levels.⁴⁸ This suggests that all models examined,

⁴⁸As discussed in Section 2.3, the null hypothesis for the test of equal forecasting performance is that $t_N = 0$, tested against the two-sided alternative that $t_N \neq 0$, where $t_N \sim \mathcal{N}(0, 1)$. Therefore, if the null hypothesis is to be rejected e.g. at the 5% or 10% significance levels, then t_N should be outside the 2.5% or 5% right- and left-tail critical values which are respectively ± 1.96 and ± 1.64 .

even the random walk benchmark, produce equally accurate 1-day-ahead density forecasts for the returns of bitcoin, ether, ripple and litecoin.

This important result extends the findings of Catania and Grassi (2021) that equal forecasting performance between an EGARCH benchmark model and more complex GAS models is the most common outcome when examining CRPS for a large number of crypto assets. Our results on the equal forecasting performance test based on the average CRPS of each model reported in Table 2.8 for bitcoin, ether, ripple and litecoin daily returns density forecasts indicate that e.g. the EGARCH model is no more accurate than the random walk benchmark for which volatility forecasts are produced with a 30-day equally-weighted moving average.

Table 2.8: Average CRPS of 1-day-ahead univariate density forecasts

	BTC	ETH	XRP	LTC
Random Walk	0.02226	0.03010	0.03436	0.03141
t EWMA(0.925)	98.92%	98.79%	98.78%	99.04%
t EWMA(0.94)	98.96%	98.87%	98.84%	99.06%
t AEWMA(0.925, 1%)	99.18%	98.87%	98.76%	99.08%
t AEWMA(0.925, 2%)	99.74%	99.07%	99.02%	99.27%
t AEWMA(0.925, 3%)	100.89%	99.50%	99.62%	99.75%
t AEWMA(0.925,-1%)	98.99%	98.75%	98.69%	99.04%
t AEWMA(0.925,-3%)	101.17%	99.52%	99.39%	99.72%
t AEWMA(0.925,-5%)	107.00%	102.19%	101.75%	102.10%
t AEWMA(0.94, 1%)	99.23%	98.96%	98.80%	99.10%
t AEWMA(0.94, 2%)	99.80%	99.17%	99.09%	99.32%
t AEWMA(0.94, 3%)	100.97%	99.62%	99.70%	99.83%
t AEWMA(0.94,-1%)	99.04%	98.83%	98.72%	99.06%
t AEWMA(0.94,-3%)	101.22%	99.61%	99.45%	99.76%
t AEWMA(0.94,-5%)	107.04%	102.27%	101.84%	102.15%
t GARCH	98.85%	97.99%	97.57%	98.79%
t EGARCH	98.70%	97.95%	97.69%	98.60%

Average CRPS of 1-day-ahead univariate density forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC) daily log returns, based on an out-of-sample period between 1 January 2017 - 31 August 2021. For the random walk benchmark model the average CRPS is reported outright and the average scores of the remaining models are expressed as a percentage of the benchmark's score. The EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE.

In the case of multivariate distribution forecasts for the returns of bitcoin, ether, ripple and litecoin, the energy score and variogram scores are calculated by drawing

10,000 random samples from the forecasted joint density of log returns produced by each model, based on the corresponding realized returns.⁴⁹ Again for comparison purposes, each score is averaged across the 4.5-year forecasting period. Table 2.9 reports the outright average scores for the random walk benchmark model and the average scores of all other models are again expressed relative to the benchmark's average score. Again, the t_N test statistic for the equal forecasting performance hypothesis test of Gneiting and Ranjan (2011) is used to determine whether the multivariate scores produced are significantly different between different models.

As shown in Table 2.9, the random walk benchmark produces a higher average energy score compared with most other models, but the lowest average variogram scores; contrary to the CRPS results, the highest multivariate average scores are produced by the ADCC- t EGARCH. However, we again find that all models exhibit equal forecasting performance: when calculating the t_N test statistic between the highest and lowest average multivariate scores, t_N is always below 0.5 for all multivariate scores, so again, the null hypothesis of equal forecasting performance is accepted e.g. at the 5% and 10% significance levels for all multivariate scores. As noted previously for CRPS, these results indicate that all models examined, even the random walk, produce equally accurate 1-day-ahead joint density forecasts for the returns of bitcoin, ether, ripple and litecoin. It therefore appears that overall, at the daily frequency for bitcoin, ether, ripple and litecoin, when volatility and covariance forecasting accuracy are evaluated via proper scoring rules, none of the multivariate or univariate parametric models attempted perform any better than the simplest 30-day EQMA random walk benchmark.

2.5.2 Hourly Forecast Evaluation

Following the discussion of out-of-sample results at the daily frequency, we now present the same analysis performed on a more recent sample of hourly bitcoin,

⁴⁹Random samples from the multivariate normal and Student- t distributions are produced using the mvnfast R package of Fasiolo (2016).

Table 2.9: Average multivariate scores for 1-day-ahead joint density forecasts

	ES	VS^{0.5}	VS¹	VS²
Random Walk	0.06727	0.04398	0.00618	0.000056
<i>t</i> EWMA(0.925)	98.47%	191.76%	336.85%	4593.03%
<i>t</i> EWMA(0.94)	98.57%	193.84%	339.93%	4604.34%
<i>t</i> AEWMA(0.925, 1%)	98.93%	197.14%	347.37%	4685.54%
<i>t</i> AEWMA(0.925, 2%)	99.18%	197.11%	347.20%	4678.14%
<i>t</i> AEWMA(0.925, 3%)	99.66%	197.15%	347.50%	4687.37%
<i>t</i> AEWMA(0.925,-1%)	98.88%	197.02%	346.98%	4677.30%
<i>t</i> AEWMA(0.925,-3%)	99.57%	197.10%	347.32%	4684.58%
<i>t</i> AEWMA(0.925,-5%)	102.07%	197.12%	347.27%	4681.18%
<i>t</i> AEWMA(0.94, 1%)	99.02%	198.55%	349.17%	4680.90%
<i>t</i> AEWMA(0.94, 2%)	99.23%	198.52%	349.15%	4680.95%
<i>t</i> AEWMA(0.94, 3%)	99.73%	198.46%	348.98%	4678.45%
<i>t</i> AEWMA(0.94,-1%)	98.90%	198.49%	348.87%	4674.88%
<i>t</i> AEWMA(0.94,-3%)	99.62%	198.47%	349.06%	4679.32%
<i>t</i> AEWMA(0.94,-5%)	102.16%	198.47%	348.90%	4676.97%
DCC- <i>t</i> GARCH	98.43%	192.49%	334.12%	4766.96%
DCC- <i>t</i> EGARCH	102.14%	270.39%	568.93%	17838.91%
ADCC- <i>t</i> GARCH	98.35%	191.65%	332.35%	4731.02%
ADCC- <i>t</i> EGARCH	102.19%	270.28%	569.76%	19283.66%

Average multivariate scores for 1-day-ahead joint density forecasts of bitcoin, ether, ripple and litecoin daily log returns, based on an out-of-sample period between 1 January 2017 - 31 August 2021. Each row reports for each model the energy score (**ES**) and variogram scores (**VS**) of order $p = 0.5, 1$ and 2 . For the random walk benchmark model each average score is reported outright and the scores of the remaining models are expressed as a percentage of the benchmark's corresponding average score. The EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE.

ether, ripple and litecoin log returns. In the following, left- and right-tail VaR and ES backtesting and also the calculation of univariate and multivariate scores are performed using hourly log returns of bitcoin, ether, ripple and litecoin and an out-of-sample period between 1 May 2021 00:00 UTC - 1 July 2021 00:00 UTC. Similar to the daily frequency analysis, we produce 1-hour-ahead volatility and covariance forecasts for bitcoin, ether, ripple and litecoin for a total of 1,465 hourly observations. For the random walk benchmark model, forecasts are based on an equally-weighted 72-hour moving average; the EWMA and AEWMA forecasts are again produced with λ set ad hoc to 0.925 and 0.94 and the AEWMA η parameter is now set to 0.7%, 0.8% and 0.9% for left-tail VaR and ES and to 0.2% and -0.2% for right-tail VaR.⁵⁰ GARCH models are calibrated using a rolling estimation window of 4 months, i.e. 2,882 hourly observations. The same distribution assumptions as in the daily frequency analysis are followed, i.e. a normal for the random walk benchmark, a Student- t distribution for the EWMA and AEWMA with ad hoc $\nu = 6$ and similar for the GARCH models with the degrees of freedom parameter estimated jointly with the model parameters via MLE.

VaR and ES Backtesting

Hourly left- and right-tail VaR backtesting results at the 1%, 2.5% and 5% significance level are shown in tables 2.10 and 2.11, where we report the total number of realized VaR exceedances x^{VaR} over the forecasting period, the probability $\Phi(z)$ of obtaining x^{VaR} or fewer realized VaR exceedances and the p-value of the conditional coverage (CC) test.

As shown in Table 2.10 for left-tail (long position) 1% hourly VaR, the AEWMA models are the only ones that produce fewer than the expected 14.65 exceedances for 1% VaR.⁵¹ The AEWMA models appear to produce accurate left-tail 1% VaR forecasts, i.e. they are in the green traffic light zone and the CC test's null hypothesis

⁵⁰Additional testing of other AEWMA η parameter choices is shown in Tables A2 and A3 of Appendix A for the 1% daily VaR forecasts and the corresponding traffic light test.

⁵¹The expected number of total VaR exceedances for 1% VaR is calculated based on equation (2.17) using $N = 1,465$ and $\alpha = 1\%$.

of an accurate VaR model is accepted even at the 10% level of significance. By comparison, it is worth noting that the left-tail 1% hourly VaR forecasts from the t EGARCH models are in the yellow zone but considered accurate based on the CC test, and the standard t GARCH forecasts are in fact not considered accurate. The right-tail (short position) hourly 1% VaR backtesting results shown in Table 2.11 appear to favour simple model specifications even more; for instance, the 1% right-tail VaR forecasts produced by the symmetric t EWMA models are considered accurate based on both the traffic light and CC tests, and even the simplest random walk benchmark model produces accurate 1% VaR forecasts for bitcoin and ripple, on par with the more complex GARCH specifications.

Similarly, hourly Expected Shortfall backtesting results are shown in Table 2.12 for left-tail (long position) ES and in Table 2.13 for the right tail (short position), reporting the realized total ES generalized exceedances x^{ES} , the probability $\Phi(z)$ of obtaining x^{ES} or fewer realized total ES generalized exceedances and the p-value of the exceedance residual (ER) test.

The hourly left-tail (long position) 1% ES backtesting results shown in Table 2.12 indicate that some of the AEWMA models produce fewer than the expected 7.325 total ES generalized exceedances, although none of the models examined achieve this across all assets examined.⁵² According to the traffic light ES test, only the forecasts based on the t AEWMA(0.94, 0.9%) specification are in the green zone for all assets examined, while the symmetric t EGARCH is ‘in the yellow’ for bitcoin and ripple. The even simpler random walk benchmark and also the symmetric t EWMA models yield ES forecasts that are in the red zone for some or all of the assets examined. Interestingly, the ES exceedance residual (ER) backtest indicates that almost all models attempted yield accurate ES forecasts even at the 10% level of rejecting the ER test’s null hypothesis of a correctly specified ES model, with the sole exception of the random walk benchmark. As noted previously for right-tail 1% VaR forecasts, the right-tail 1% ES forecasts shown in Table 2.13 again favour simpler models, and

⁵²The expected value of the total ES generalized exceedances for 1% ES is calculated based on equation (2.22) using $N = 1,465$ and $\alpha = 1\%$.

forecasts produced even from the symmetric t EWMA model are on par with the more complex GARCH specifications, with the random walk benchmark again being the sole exception that produces inaccurate forecasts.

Moreover, as in the case of daily data, the 2.5% and the 5% left- and right-tail VaR and ES backtesting results shown in tables 2.10, 2.11, 2.12 and 2.13 also indicate that EWMA-type models, and sometimes even the random walk benchmark, can produce accurate forecasts at higher VaR and ES significance levels and are therefore again on par with the more complex GARCH models.

Distribution Forecast Evaluation

The hourly forecast distribution evaluation results are shown in Table 2.14 for CRPS and in Table 2.15 for the multivariate scores. Similar to the daily forecast evaluation shown previously, the random walk benchmark yields a higher average CPRS in the univariate forecasting case compared with most other models and in the multivariate case it exhibits the highest average energy score but the lowest average variogram scores; a t EGARCH specification yields the lowest average CRPS for assets examined except ripple for which t GARCH average CRPS is slightly lower, and similarly the ADCC- t GARCH yields the lowest average energy score. Importantly and similar to the discussion presented in the daily frequency analysis, the test of equal forecasting performance does not distinguish between the models: the t_N test statistic between the highest and lowest average scores is always below 0.45, so the null hypothesis of equal forecasting performance is accepted e.g. at the 5% and 10% significance levels for all assets examined and for all univariate and multivariate scores.⁵³ This suggests that the equal forecasting performance identified previously based on the average scores of daily returns density forecasts also holds at the hourly frequency: i.e. all models examined, even the random walk benchmark, produce equally accurate

⁵³As discussed in Section 2.3, the null hypothesis for the test of equal forecasting performance is that $t_N = 0$, tested against the two-sided alternative that $t_N \neq 0$, where $t_N \sim \mathcal{N}(0, 1)$. Therefore, if the null hypothesis is to be rejected e.g. at the 5% or 10% significance levels, then t_N should be outside the 2.5% or 5% right- and left-tail critical values which are respectively ± 1.96 and ± 1.64 .

Table 2.10: Backtesting results for 1-hour-ahead left-tail VaR

Hourly 1% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	36	1.0000	0.000***	36	1.0000	0.000***	38	1.0000	0.000***	36	1.0000	0.000***
t EWMA(0.925)	29	0.9999	0.001***	29	0.9999	0.001***	29	0.9999	0.003***	26	0.9986	0.017**
t EWMA(0.94)	28	0.9998	0.007***	27	0.9994	0.004***	27	0.9994	0.012**	24	0.9930	0.054*
t AEWMA(0.925, 0.7%)	15	0.5366	0.340	15	0.5366	0.853	21	0.9523	0.174	14	0.4322	0.861
t AEWMA(0.925, 0.8%)	11	0.1689	0.557	14	0.4322	0.861	19	0.8733	0.280	13	0.3324	0.807
t AEWMA(0.925, 0.9%)	9	0.0690	0.264	13	0.3324	0.807	18	0.8105	0.325	13	0.3324	0.807
t AEWMA(0.94, 0.7%)	15	0.5366	0.340	15	0.5366	0.853	20	0.9200	0.227	14	0.4322	0.861
t AEWMA(0.94, 0.8%)	9	0.0690	0.264	14	0.4322	0.861	18	0.8105	0.325	14	0.4322	0.861
t AEWMA(0.94, 0.9%)	8	0.0404	0.154	11	0.1689	0.557	18	0.8105	0.325	12	0.2433	0.700
t GARCH	23	0.9858	0.087*	19	0.8733	0.429	26	0.9986	0.021**	24	0.9930	0.054*
t EGARCH	21	0.9523	0.174	19	0.8733	0.429	22	0.9732	0.126	21	0.9523	0.216

Hourly 2.5% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	55	0.9989	0.016**	55	0.9989	0.000***	58	0.9998	0.002***	54	0.9982	0.001***
t EWMA(0.925)	59	0.9999	0.002***	52	0.9950	0.019**	54	0.9982	0.020**	52	0.9950	0.002***
t EWMA(0.94)	57	0.9997	0.004***	51	0.9919	0.024**	52	0.9950	0.038**	50	0.9874	0.001***
t AEWMA(0.925, 0.7%)	29	0.1010	0.364	38	0.5910	0.057*	35	0.3928	0.948	35	0.3928	0.409
t AEWMA(0.925, 0.8%)	26	0.0377	0.135	31	0.1733	0.578	34	0.3302	0.882	29	0.1010	0.232
t AEWMA(0.925, 0.9%)	19	0.0016	0.003***	29	0.1010	0.364	29	0.1010	0.364	28	0.0745	0.187
t AEWMA(0.94, 0.7%)	28	0.0745	0.272	36	0.4584	0.186	35	0.3928	0.948	34	0.3302	0.404
t AEWMA(0.94, 0.8%)	25	0.0259	0.089*	31	0.1733	0.578	32	0.2195	0.689	29	0.1010	0.232
t AEWMA(0.94, 0.9%)	21	0.0045	0.011**	28	0.0745	0.272	29	0.1010	0.364	25	0.0259	0.077*
t GARCH	53	0.9969	0.037**	50	0.9874	0.031**	45	0.9195	0.376	44	0.8914	0.125
t EGARCH	48	0.9715	0.169	47	0.9587	0.134	42	0.8158	0.666	41	0.7680	0.237

Hourly 5% VaR Backtesting – Long position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	82	0.8529	0.179	81	0.8236	0.086*	86	0.9368	0.142	82	0.8529	0.087*
t EWMA(0.925)	89	0.9705	0.147	98	0.9985	0.000***	96	0.9968	0.003***	96	0.9968	0.001***
t EWMA(0.94)	86	0.9368	0.330	94	0.9936	0.028**	93	0.9910	0.009***	97	0.9978	0.000***
t AEWMA(0.925, 0.7%)	56	0.0193	0.098*	64	0.1337	0.042**	69	0.3052	0.802	67	0.2269	0.092*
t AEWMA(0.925, 0.8%)	49	0.0018	0.008***	59	0.0438	0.008***	59	0.0438	0.194	61	0.0710	0.016**
t AEWMA(0.925, 0.9%)	45	0.0004	0.001***	53	0.0076	0.000***	53	0.0076	0.029**	54	0.0105	0.009***
t AEWMA(0.94, 0.7%)	53	0.0076	0.039**	66	0.1924	0.073*	67	0.2269	0.749	68	0.2646	0.042**
t AEWMA(0.94, 0.8%)	48	0.0012	0.005***	59	0.0438	0.008***	61	0.0710	0.306	60	0.0561	0.034**
t AEWMA(0.94, 0.9%)	45	0.0004	0.001***	53	0.0076	0.000***	55	0.0143	0.060*	55	0.0143	0.014**
t GARCH	85	0.9205	0.344	89	0.9705	0.101	93	0.9910	0.051*	87	0.9504	0.035**
t EGARCH	89	0.9705	0.184	83	0.8788	0.087*	93	0.9910	0.051*	79	0.7547	0.786

Backtesting results for 1-hour-ahead left-tail 1%, 2.5% and 5% VaR forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 May 2021 - 1 July 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the total number of realized VaR exceedances x^{VaR} ; the probability $\Phi(z)$ of obtaining the realized number of VaR exceedances; the p-value of the CC test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the CC test, the null hypothesis is that the VaR forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

Table 2.11: Backtesting results for 1-hour-ahead right-tail VaR

Hourly 1% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	19	0.8733	0.280	27	0.9994	0.009***	20	0.9200	0.313	24	0.9930	0.054*
t EWMA(0.925)	16	0.6385	0.788	16	0.6385	0.788	20	0.9200	0.227	16	0.6385	0.788
t EWMA(0.94)	15	0.5366	0.853	15	0.5366	0.853	17	0.7314	0.683	16	0.6385	0.788
t AEWMA(0.925, 0.2%)	17	0.7314	0.683	17	0.7314	0.683	17	0.7314	0.683	16	0.6385	0.788
t AEWMA(0.925, -0.2%)	15	0.5366	0.853	16	0.6385	0.788	18	0.8105	0.557	17	0.7314	0.683
t AEWMA(0.94, 0.2%)	12	0.2433	0.700	17	0.7314	0.683	17	0.7314	0.683	16	0.6385	0.788
t AEWMA(0.94, -0.2%)	15	0.5366	0.853	15	0.5366	0.853	16	0.6385	0.788	16	0.6385	0.788
t GARCH	13	0.3324	0.807	12	0.2433	0.700	15	0.5366	0.853	13	0.3324	0.807
t EGARCH	12	0.2433	0.700	12	0.2433	0.700	13	0.3324	0.807	12	0.2433	0.700

Hourly 2.5% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	37	0.5250	0.613	37	0.5250	0.382	38	0.5910	0.635	33	0.2721	0.382
t EWMA(0.925)	29	0.1010	0.364	33	0.2721	0.793	39	0.6545	0.246	30	0.1338	0.467
t EWMA(0.94)	29	0.1010	0.133	32	0.2195	0.689	40	0.7139	0.253	30	0.1338	0.467
t AEWMA(0.925, 0.2%)	26	0.0377	0.108	32	0.2195	0.689	36	0.4584	0.186	28	0.0745	0.272
t AEWMA(0.925, -0.2%)	29	0.1010	0.133	30	0.1338	0.467	40	0.7139	0.253	27	0.0536	0.145
t AEWMA(0.94, 0.2%)	25	0.0259	0.077*	32	0.2195	0.689	39	0.6545	0.637	28	0.0745	0.272
t AEWMA(0.94, -0.2%)	29	0.1010	0.364	31	0.1733	0.578	35	0.3928	0.518	27	0.0536	0.145
t GARCH	28	0.0745	0.272	30	0.1338	0.467	32	0.2195	0.689	25	0.0259	0.077*
t EGARCH	25	0.0259	0.077*	29	0.1010	0.232	29	0.1010	0.232	26	0.0377	0.108

Hourly 5% VaR Backtesting – Short position												
	x^{VaR}	BTC		x^{VaR}	ETH		x^{VaR}	XRP		x^{VaR}	LTC	
		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC		$\Phi(z)$	CC
Random Walk	55	0.0143	0.014**	71	0.3937	0.919	62	0.0887	0.149	49	0.0018	0.008***
t EWMA(0.925)	63	0.1096	0.338	72	0.4404	0.958	62	0.0887	0.023**	67	0.2269	0.032**
t EWMA(0.94)	61	0.0710	0.216	75	0.5831	0.818	60	0.0561	0.034**	61	0.0710	0.114
t AEWMA(0.925, 0.2%)	56	0.0193	0.083*	70	0.3484	0.866	61	0.0710	0.047**	63	0.1096	0.031**
t AEWMA(0.925, -0.2%)	61	0.0710	0.306	76	0.6292	0.827	62	0.0887	0.023**	64	0.1337	0.112
t AEWMA(0.94, 0.2%)	56	0.0193	0.049**	66	0.1924	0.677	59	0.0438	0.062*	57	0.0257	0.031**
t AEWMA(0.94, -0.2%)	61	0.0710	0.306	74	0.5358	0.585	61	0.0710	0.047**	63	0.1096	0.192
t GARCH	59	0.0438	0.194	77	0.6735	0.451	60	0.0561	0.085*	55	0.0143	0.051*
t EGARCH	62	0.0887	0.373	71	0.3937	0.654	55	0.0143	0.005***	54	0.0105	0.042**

Backtesting results for 1-hour-ahead right-tail 1%, 2.5% and 5% VaR forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 May 2021 - 1 July 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the total number of realized VaR exceedances x^{VaR} ; the probability $\Phi(z)$ of obtaining the realized number of VaR exceedances; the p-value of the CC test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the CC test, the null hypothesis is that the VaR forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

Table 2.12: Backtesting results for 1-hour-ahead left-tail ES

Hourly 1% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	25.5	1.0000	0.00***	24.0	1.0000	0.02**	29.0	1.0000	0.00***	26.8	1.0000	0.00***
t EWMA(0.925)	19.3	1.0000	0.02**	13.5	0.9974	0.40	19.6	1.0000	0.02**	15.0	0.9998	0.29
t EWMA(0.94)	18.1	1.0000	0.04**	12.5	0.9904	0.37	18.1	1.0000	0.03**	14.6	0.9995	0.23
t AEWMA(0.925, 0.7%)	7.7	0.5745	0.26	8.0	0.6237	0.32	13.2	0.9961	0.11	8.5	0.7087	0.25
t AEWMA(0.925, 0.8%)	6.2	0.3000	0.21	7.1	0.4559	0.32	12.1	0.9848	0.13	8.0	0.6147	0.27
t AEWMA(0.925, 0.9%)	5.2	0.1713	0.21	6.1	0.2866	0.36	11.1	0.9575	0.13	7.3	0.4978	0.33
t AEWMA(0.94, 0.7%)	7.4	0.5201	0.33	7.4	0.5100	0.31	12.3	0.9876	0.14	8.5	0.7067	0.25
t AEWMA(0.94, 0.8%)	5.8	0.2419	0.14	6.4	0.3363	0.36	11.4	0.9681	0.14	7.8	0.5894	0.34
t AEWMA(0.94, 0.9%)	5.1	0.1589	0.16	5.6	0.2157	0.32	10.5	0.9225	0.21	7.2	0.4755	0.29
t GARCH	13.6	0.9977	0.24	9.7	0.8587	0.34	15.1	0.9998	0.28	12.4	0.9897	0.49
t EGARCH	12.1	0.9849	0.29	10.0	0.8881	0.36	13.4	0.9971	0.36	10.7	0.9401	0.40

Hourly 2.5% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	37.6	1.0000	0.00***	37.6	1.0000	0.00***	41.1	1.0000	0.00***	37.6	1.0000	0.00***
t EWMA(0.925)	35.0	1.0000	0.05**	29.6	0.9994	0.12	33.1	1.0000	0.01**	30.2	0.9997	0.10
t EWMA(0.94)	33.4	1.0000	0.06*	28.4	0.9983	0.15	31.9	1.0000	0.04**	29.3	0.9992	0.11
t AEWMA(0.925, 0.7%)	16.3	0.2852	0.20	18.3	0.4942	0.35	22.0	0.8542	0.02**	17.9	0.4568	0.29
t AEWMA(0.925, 0.8%)	13.7	0.0918	0.27	16.3	0.2760	0.32	20.1	0.7014	0.04**	16.3	0.2787	0.26
t AEWMA(0.925, 0.9%)	11.3	0.0214	0.21	14.5	0.1365	0.37	18.5	0.5244	0.03**	14.9	0.1588	0.34
t AEWMA(0.94, 0.7%)	15.8	0.2295	0.22	17.6	0.4190	0.37	21.2	0.7987	0.05**	18.0	0.4662	0.28
t AEWMA(0.94, 0.8%)	13.3	0.0739	0.27	15.6	0.2204	0.34	19.5	0.6289	0.04**	16.4	0.2929	0.23
t AEWMA(0.94, 0.9%)	11.1	0.0185	0.32	13.9	0.0997	0.37	17.9	0.4528	0.05*	15.1	0.1731	0.19
t GARCH	28.4	0.9982	0.32	25.1	0.9749	0.39	26.0	0.9871	0.11	24.4	0.9616	0.18
t EGARCH	26.7	0.9923	0.28	23.4	0.9284	0.41	25.2	0.9768	0.13	23.6	0.9369	0.16

Hourly 5% ES Backtesting – Long position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	52.3	0.9994	0.00***	52.9	0.9996	0.00***	56.8	1.0000	0.00***	52.8	0.9996	0.00***
t EWMA(0.925)	54.4	0.9999	0.01**	51.5	0.9990	0.20	54.8	0.9999	0.04**	52.8	0.9996	0.14
t EWMA(0.94)	53.3	0.9997	0.01***	49.3	0.9957	0.18	53.0	0.9996	0.03**	51.0	0.9985	0.17
t AEWMA(0.925, 0.7%)	30.4	0.0980	0.21	34.2	0.3054	0.26	35.6	0.4161	0.13	33.9	0.2869	0.31
t AEWMA(0.925, 0.8%)	26.5	0.0185	0.26	31.0	0.1248	0.31	32.8	0.2121	0.09*	30.6	0.1080	0.35
t AEWMA(0.925, 0.9%)	22.8	0.0022	0.39	28.3	0.0437	0.35	30.3	0.0946	0.09*	27.7	0.0323	0.35
t AEWMA(0.94, 0.7%)	29.7	0.0777	0.18	33.8	0.2789	0.34	35.4	0.4032	0.12	33.2	0.2392	0.35
t AEWMA(0.94, 0.8%)	26.1	0.0152	0.26	30.7	0.1098	0.32	32.6	0.2055	0.13	29.9	0.0841	0.33
t AEWMA(0.94, 0.9%)	22.6	0.0019	0.39	28.0	0.0371	0.34	30.1	0.0894	0.12	27.1	0.0246	0.33
t GARCH	50.0	0.9971	0.08*	47.0	0.9841	0.32	49.3	0.9955	0.27	44.2	0.9398	0.30
t EGARCH	49.2	0.9953	0.17	45.1	0.9601	0.22	46.5	0.9788	0.44	41.1	0.8228	0.23

Backtesting results for 1-hour-ahead left-tail 1%, 2.5% and 5% ES forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 May 2021 - 1 July 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the realized total ES generalized exceedances x^{ES} ; the probability $\Phi(z)$ of obtaining the realized total ES generalized exceedances; the p-value of the ER test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the ER test, the null hypothesis is that the ES forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

Table 2.13: Backtesting results for 1-hour-ahead right-tail ES

Hourly 1% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	12.7	0.9931	0.03**	17.5	1.0000	0.02**	15.5	0.9999	0.00***	17.8	1.0000	0.00***
t EWMA(0.925)	9.2	0.8055	0.34	9.3	0.8202	0.34	10.4	0.9182	0.33	9.7	0.8556	0.48
t EWMA(0.94)	8.4	0.6939	0.45	8.8	0.7503	0.33	9.3	0.8124	0.38	9.2	0.8069	0.54
t AEWMA(0.925, 0.2%)	8.3	0.6784	0.51	8.8	0.7441	0.48	9.5	0.8401	0.28	9.5	0.8423	0.53
t AEWMA(0.925,-0.2%)	8.7	0.7279	0.41	9.0	0.7721	0.37	9.8	0.8654	0.32	8.8	0.7425	0.65
t AEWMA(0.94, 0.2%)	7.5	0.5357	0.28	8.3	0.6684	0.52	8.6	0.7198	0.42	9.1	0.7866	0.59
t AEWMA(0.94,-0.2%)	8.0	0.6208	0.59	8.4	0.6874	0.35	8.9	0.7697	0.38	8.3	0.6749	0.64
t GARCH	5.2	0.1725	0.90	5.5	0.2016	0.65	6.9	0.4247	0.63	6.5	0.3559	0.80
t EGARCH	5.7	0.2282	0.62	5.4	0.1853	0.70	5.8	0.2510	0.69	7.4	0.5095	0.48

Hourly 2.5% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	22.5	0.8882	0.04**	25.7	0.9838	0.00***	23.5	0.9347	0.01***	24.5	0.9629	0.00***
t EWMA(0.925)	18.6	0.5287	0.11	18.7	0.5401	0.26	21.7	0.8361	0.32	17.1	0.3589	0.19
t EWMA(0.94)	17.6	0.4173	0.21	18.4	0.5061	0.23	20.1	0.7001	0.44	16.8	0.3293	0.32
t AEWMA(0.925, 0.2%)	17.0	0.3564	0.12	18.3	0.5018	0.26	20.9	0.7686	0.19	16.7	0.3195	0.16
t AEWMA(0.925,-0.2%)	17.8	0.4388	0.20	17.2	0.3770	0.18	20.6	0.7465	0.41	16.7	0.3166	0.15
t AEWMA(0.94, 0.2%)	16.3	0.2767	0.15	18.0	0.4626	0.31	18.9	0.5725	0.49	16.4	0.2868	0.23
t AEWMA(0.94,-0.2%)	16.9	0.3422	0.30	16.8	0.3295	0.30	19.1	0.5930	0.34	16.4	0.2855	0.24
t GARCH	14.1	0.1120	0.82	14.5	0.1327	0.73	16.5	0.3046	0.44	13.7	0.0917	0.51
t EGARCH	14.2	0.1161	0.57	14.0	0.1088	0.71	15.0	0.1662	0.50	13.8	0.0977	0.47

Hourly 5% ES Backtesting – Short position												
	x^{ES}	BTC		x^{ES}	ETH		x^{ES}	XRP		x^{ES}	LTC	
		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER		$\Phi(z)$	ER
Random Walk	34.0	0.2973	0.01***	40.9	0.8090	0.03**	36.4	0.4816	0.01**	32.2	0.1800	0.00***
t EWMA(0.925)	32.4	0.1891	0.49	36.3	0.4735	0.53	37.5	0.5752	0.12	32.8	0.2150	0.56
t EWMA(0.94)	31.8	0.1609	0.44	34.8	0.3539	0.69	36.2	0.4663	0.14	31.1	0.1281	0.42
t AEWMA(0.925, 0.2%)	29.6	0.0730	0.39	34.8	0.3516	0.58	35.7	0.4257	0.18	30.9	0.1179	0.57
t AEWMA(0.925,-0.2%)	31.3	0.1348	0.51	34.8	0.3569	0.71	37.1	0.5373	0.13	31.9	0.1671	0.51
t AEWMA(0.94, 0.2%)	28.8	0.0538	0.42	33.2	0.2407	0.55	34.5	0.3286	0.19	29.3	0.0654	0.45
t AEWMA(0.94,-0.2%)	30.8	0.1137	0.49	33.9	0.2847	0.72	35.9	0.4409	0.18	30.4	0.1002	0.55
t GARCH	29.1	0.0608	0.77	33.1	0.2356	0.97	31.9	0.1669	0.54	26.7	0.0206	0.68
t EGARCH	27.4	0.0279	0.85	33.0	0.2257	0.97	28.7	0.0511	0.63	27.6	0.0309	0.61

Backtesting results for 1-hour-ahead right-tail 1%, 2.5% and 5% ES forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC), based on an out-of-sample period between 1 May 2021 - 1 July 2021. For each asset, the first column denotes the models used, where the EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE. The remaining columns denote: the realized total ES generalized exceedances x^{ES} ; the probability $\Phi(z)$ of obtaining the realized total ES generalized exceedances; the p-value of the ER test. For $\Phi(z)$, the traffic light zones are defined as: **Green** if the probability is below 0.95, **Yellow** if it is greater (or equal) than 0.95 and less than 0.9999; and **Red** if it exceeds (or is equal to) 0.9999. For the ER test, the null hypothesis is that the ES forecasting model is specified correctly and the p-values are denoted with *, **, *** if the null hypothesis is rejected at the 10%, 5% and 1% significance level respectively.

1-hour-ahead univariate and multivariate density forecasts for the returns of bitcoin, ether, ripple and litecoin.

Table 2.14: Average CRPS of 1-hour-ahead univariate density forecasts

	BTC	ETH	XRP	LTC
Random Walk	0.00613	0.00858	0.00920	0.00890
t EWMA(0.925)	99.25%	99.13%	98.44%	99.25%
t EWMA(0.94)	99.19%	99.11%	98.43%	99.22%
t AEWMA(0.925, 0.2%)	99.15%	99.04%	98.40%	99.16%
t AEWMA(0.925, 0.7%)	99.85%	99.01%	98.70%	99.16%
t AEWMA(0.925, 0.8%)	100.32%	99.12%	98.89%	99.28%
t AEWMA(0.925, 0.9%)	100.94%	99.30%	99.13%	99.46%
t AEWMA(0.925,-0.2%)	99.36%	99.20%	98.49%	99.32%
t AEWMA(0.94, 0.2%)	99.10%	99.02%	98.40%	99.14%
t AEWMA(0.94, 0.7%)	99.87%	99.04%	98.74%	99.20%
t AEWMA(0.94, 0.8%)	100.35%	99.17%	98.94%	99.33%
t AEWMA(0.94, 0.9%)	100.97%	99.35%	99.19%	99.52%
t AEWMA(0.94,-0.2%)	99.30%	99.17%	98.49%	99.28%
t GARCH	98.91%	98.82%	98.04%	98.89%
t EGARCH	98.56%	98.63%	98.10%	98.73%

Average CRPS of 1-hour-ahead univariate density forecasts for bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC) hourly log returns, based on an out-of-sample period between 1 May 2021 - 1 July 2021. For the random walk benchmark model the average CRPS is reported outright and the average scores of the remaining models are expressed as a percentage of the benchmark's score. The EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE.

2.6 Summary and Conclusions

In this chapter, we conduct an out-of-sample analysis on bitcoin, ether, ripple and litecoin using both daily and hourly log returns. Daily forecasts are produced between January 2017 - August 2021 and hourly forecasts between 1 May 2021 - 1 July 2021. Volatility is modelled using a simple equally-weighted moving average benchmark model (random walk) and also RiskMetrics-type EWMA models, an asymmetric EWMA specification and GARCH models. In order to produce multivariate density forecasts, the correlation structure is modelled via the multivariate random walk and EWMA models and also using the symmetric and asymmetric DCC models. The

Table 2.15: Average multivariate scores for 1-hour-ahead joint density forecasts

	ES	VS^{0.5}	VS¹	VS²
Random Walk	0.01788	0.00519	0.00008	0.00000002
<i>t</i> EWMA(0.925)	98.67%	305.95%	895.48%	22178.19%
<i>t</i> EWMA(0.94)	98.69%	307.40%	900.18%	22158.81%
<i>t</i> AEWMA(0.925, 0.2%)	98.90%	314.23%	936.95%	23342.87%
<i>t</i> AEWMA(0.925, 0.7%)	99.06%	314.31%	936.87%	23342.93%
<i>t</i> AEWMA(0.925, 0.8%)	99.27%	314.33%	936.83%	23355.64%
<i>t</i> AEWMA(0.925, 0.9%)	99.46%	314.38%	937.20%	23350.47%
<i>t</i> AEWMA(0.925,-0.2%)	99.05%	314.13%	936.26%	23334.51%
<i>t</i> AEWMA(0.94, 0.2%)	98.88%	314.76%	936.31%	23135.99%
<i>t</i> AEWMA(0.94, 0.7%)	99.14%	314.68%	936.27%	23136.28%
<i>t</i> AEWMA(0.94, 0.8%)	99.21%	314.86%	936.89%	23144.19%
<i>t</i> AEWMA(0.94, 0.9%)	99.50%	314.75%	936.33%	23139.95%
<i>t</i> AEWMA(0.94,-0.2%)	99.03%	314.73%	936.40%	23144.35%
DCC- <i>t</i> GARCH	98.71%	312.46%	911.37%	21879.95%
DCC- <i>t</i> EGARCH	98.60%	313.16%	905.05%	21706.41%
ADCC- <i>t</i> GARCH	98.68%	311.88%	907.76%	21739.47%
ADCC- <i>t</i> EGARCH	98.62%	312.99%	902.99%	21720.21%

Average multivariate scores for 1-hour-ahead joint density forecasts of bitcoin, ether, ripple and litecoin hourly log returns, based on an out-of-sample period between 1 May 2021 - 1 July 2021. Each row reports for each model the energy score (**ES**) and variogram scores (**VS**) of order $p = 0.5$, 1 and 2. For the random walk benchmark model each average score is reported outright and the scores of the remaining models are expressed as a percentage of the benchmark's corresponding average score. The EWMA, AEWMA and GARCH models are based on a Student- t distribution assumption and the degrees of freedom for EWMA and AEWMA are set to the ad hoc $\nu = 6$ and for GARCH the degrees of freedom are estimated via MLE.

1-period-ahead left- and right-tail Value-at-Risk and Expected Shortfall are forecasted for each asset and at multiple significance levels. VaR forecasts are backtested using the conditional coverage and the industry standard traffic light tests and similarly, ES is backtested with the exceedance residual test and again a modified traffic light test. Returns distribution forecasts are also assessed using the continuous ranked probability score (CRPS) for each asset and multivariate scoring rules such as the energy and variogram scores for the joint density forecasts of bitcoin, ether, ripple and litecoin returns.

Overall, the daily left- and right-tail VaR and ES backtesting results are consistent with the relevant literature as discussed in Section 2.2, such as Liu et al. (2020), in that daily VaR and ES forecasting models are more accurate as the VaR and ES significance level increases. Moreover, the daily backtesting results shown in Section 2.5 extend the findings of Catania and Grassi (2021) to even simpler models compared with the EGARCH benchmark used by the authors. Catania and Grassi (2021) find that 1% daily left-tail VaR forecasts produced by an EGARCH benchmark model are often on par with more complex GAS model specifications. Our results for both daily and hourly VaR and ES forecasts for both the left and right tail show that the asymmetric EWMA with ad hoc parameter choices are almost always on par with a more complex EGARCH specification. This finding sometimes extends to the simpler RiskMetrics-type symmetric EWMA specifications and even to the equally-weighted moving average random walk model.

Regarding the evaluation of the returns density forecasts for bitcoin, ether, ripple and litecoin, and the test of equal forecasting performance, all models examined, even the random walk benchmark, are found to be equally accurate both at the univariate and multivariate forecasting case and both at the daily and hourly frequencies. This important result extends the findings of Catania and Grassi (2021) on crypto asset returns distribution forecasting and indicates that when producing 1-period-ahead forecasts for the univariate or multivariate daily or hourly returns distributions of bitcoin, ether, ripple and litecoin, an EGARCH model with a Student- t distribution

assumption and e.g. an ADCC covariance model may be just as accurate as the simple equally-weighted moving average volatility and covariance benchmark model.

BLOCKCHAIN TRANSACTIONS AND CRYPTO ASSET MARKET MANIPULATION

3.1	Introduction	85
3.2	Literature Review	87
3.3	Methodology	99
3.4	Data	105
3.4.1	Market Data	106
3.4.2	Blockchain Data	109
3.4.3	Tether on the Omni Layer	116
3.5	Empirical Results	126
3.6	Summary and Conclusions	130

3.1 Introduction

This Chapter examines the topic of crypto asset market manipulation via tether blockchain transactions, developing a methodology derived from securities fraud litigation studies to identify blockchain transactions with a potentially manipulative effect on crypto asset prices. Large blockchain transactions of the tether (USDT) stablecoin originating from the (centralized) tether treasury are examined and assessed as to whether they produce significant positive abnormal returns for bitcoin prices, i.e. a material price impact. The methodology is adapted from single-firm event studies used in securities fraud litigation, using regression factor models. The results produced from this methodology can be useful in determining materiality and estimating potential damages in legal cases of crypto asset market manipulation.

Tether (USDT) is the first and highest-cap stablecoin – i.e. fixed-price crypto asset, established in 2015. Twomey and Mann (2020) provide an overview of tether’s history until the end of 2018. They discuss its controversial banking and audit history, its proven relationship with the Bitfinex crypto asset exchange, the allegations that its 1-to-1 dollar peg has not always been backed by adequate dollar reserves and the further allegations that non-collateralized tether tokens were used to artificially inflate crypto asset prices or at least prevent them from dropping. All of these allegations can be traced back to blog posts in 2017 and an online report providing some analysis of tether’s effect on bitcoin prices.¹ There followed other attempts to examine tether’s on-chain activity such as the post of Vicati (2018) on Hacker Noon and the paper of Griffin and Shams (2020) who first published their findings on SSRN in June 2018.

Several interesting tether-related events took place since the beginning of 2019. In April 2019 the New York Attorney General (NYAG) initiated a court case against Bitfinex alleging that it had ‘borrowed’ \$850 million of tether’s dollar reserve to cover losses; this was followed by an admission by Bitfinex’s legal representatives that

¹See e.g. the blog post on Medium by user [Bitfinex’ed](#) and a [discussion](#) of the online tether report.

only 74% of tether’s supply was backed by USD reserves.² Around the same time, tether’s supply began to increase from approximately 2 billion USDT in late March to 4 billion by mid-July 2019, to approximately 20 billion by early 2021, and even surpassing 80 billion temporarily in mid-2022. The outcome of the NYAG case was an out-of-court settlement in February 2021, in which Bitfinex and Tether – both subsidiaries of the same Hong Kong-registered company iFinex, paid \$18.5 million in fines and agreed to cease USDT trading in New York, while not admitting to any explicit wrongdoing.

Later in 2021, 49% of tether’s dollar reserves are revealed to be ‘unspecified commercial paper’ (Protos, 2021*b*) and large amounts of newly-issued tether supply were found to have been sent from the tether treasury to large crypto asset market makers (Protos, 2021*a*). As tether-related legal cases are still ongoing and the extant literature on stablecoins does not provide a simple and direct methodology that could be applicable to such legal cases, this Chapter provides a methodology that identifies stablecoin blockchain transactions with a potentially manipulative effect on crypto asset prices.

In the following, Section 3.2 presents an overview of the relevant literature on stablecoins and their effect on crypto asset prices and valuation; in Section 3.3 the regression factor models used to estimate abnormal bitcoin returns and the hypothesis test methodology to determine their significance are discussed; Section 3.4 provides an overview of the hourly market and blockchain transaction historical data, also providing a detailed description of the data retrieval process from a blockchain node, based on the experience of setting up and operating a blockchain node locally; Section 3.5 presents the empirical results of the regression factor models, focussing on material USDT outflow events from the tether treasury towards crypto asset exchanges and other unknown entities, i.e. events that produce significant positive

²Yet more cases were filed against tether and Bitfinex, even involving other crypto asset exchanges such as Bittrex and Poloniex, as potential conduits in the alleged Bitfinex-instigated scheme to manipulate crypto asset prices using tether.

abnormal bitcoin returns; finally, Section 3.6 provides a summary and concluding remarks.

3.2 Literature Review

The effect of stablecoins on crypto assets is an emerging topic in the crypto asset financial literature; a relevant Scopus search yields 46 papers published between 2018 – mid-2022 in Economics, Econometrics, Finance, Business, Management or Accounting journals and/or conferences, more than 90% of which are published in or after 2020.³ Stablecoin finance is examined in a variety of contexts. For instance, Hoang and Baur (2021) examine their price stability, their role in trading facilitation and their safe-haven properties in the crypto asset class, while Wang et al. (2020) study stablecoin properties as safe-havens, diversifiers and/or hedges against traditional financial assets. Lyons and Viswanath-Natraj (2020) also explore inter alia the dynamics of stablecoin price pegs, and Barucci et al. (2021) study stablecoins as a store-of-value, beyond their role as safe-havens and in crypto asset trading facilitation. Smith (2021) focuses on the use of stablecoins in DeFi lending platforms and Hampl and Gyönyörová (2021) examine stablecoin risk characteristics in relation to their potential characterization as cash-equivalent instruments under international accounting standards.

However, an emerging stream in stablecoin-relevant financial research, which this Chapter focuses on, studies the effect of stablecoins on crypto asset prices. Table 3.1 presents a summary of the most relevant papers in this stream, their research questions, methodologies, assets examined and data and sample periods used.

As shown in Table 3.1, all relevant papers examine tether (USDT) and several examine multiple stablecoins, which includes USD Coin (USDC), Paxos Standard (PAX), Dai (DAI), Binance USD (BUSD), TrueUSD (TUSD), Huobi USD (HUSD),

³These results are produced with the following Scopus search query: TITLE-ABS-KEY (("tether" OR "stablecoin" OR "stable coin") AND ("cryptocurrency" OR "cryptoasset" OR "crypto asset")) AND (LIMIT-TO (SUBJAREA , "ECON") OR LIMIT-TO (SUBJAREA , "BUSI")) .

Table 3.1: Literature of the effect of stablecoins on crypto assets

Authors	Assets	Research Questions	Methodology	Data	Sample period
Kristoufek (2022)	28 stablecoins 14 crypto assets	Does the amount of stablecoins in circulation reflect demand factors in the market? Is there evidence of suspicious bubble-boosting dynamics?	VAR VECM Causality tests	Daily market cap	2015-2020
Kristoufek (2021)	10 stablecoins BTC, ETH XRP	Examine the interaction between stablecoins and major crypto assets	Directed spillovers VAR VECM	Daily crypto asset prices and stablecoin supply	2016-2021
Ante et al. (2021a)	6 stablecoins BTC, ETH XRP, LTC	Do large stablecoin transfers affect Bitcoin returns and volume?	Event study	Hourly crypto asset returns Stablecoin blockchain transactions	2019-2020
Ante et al. (2021b)	7 stablecoins BTC, ETH XRP, LTC	Analyze crypto asset returns before and after issuance of stablecoins	Event study	Hourly crypto asset returns Stablecoin blockchain transactions	2019-2020
Grobys and Huynh (2021)	USDT BTC	Are jumps in USDT returns Granger-causal for BTC returns?	VAR Asymptotic distribution theory	Hourly BTC/USD & USDT/USD returns	2018-2021
Griffin and Shams (2020)	USDT 14 crypto assets	Did USDT influence crypto asset prices in 2017?	Linear regression Event study	Hourly crypto asset returns BTC & USDT blockchain transactions	2017-2018
Lyons and Viswanath-Natraj (2020)	USDT BTC, ETH	Does USDT play a role via intervention/ manipulation in currency markets?	VAR Impulse response functions	BTC/USD & ETH/USD daily returns USDT blockchain transactions	2017-2020
Wei (2018)	USDT BTC	Examine the impact of USDT issuance on BTC valuation	VAR Granger causality test	BTC/USD daily returns USDT blockchain transactions	2016-2018

Key characteristics of the relevant academic papers that examine the effect of stablecoins on crypto assets. The columns indicate for each paper the author(s), the crypto assets examined, the research questions, the methodologies used, the type of data and sample period employed.

Stasis Euro (EURS), QCash (QC), USDK, sUSD, Neutrino Dollar (USDN), Just (JST), Gemini Dollar (GUSD), 1SG, Anchor (ANCT), USDQ, CryptoFrank (XCHF), VNDC, USDJ, bitCNY, EURBase, EOSDT, Constant (CONST), USDx, bitUSD, NuBits and Egoras Dollar. Similarly, the most common crypto assets examined are bitcoin (BTC), ether (ETH) and ripple (XRP), with other assets included in some cases, such as litecoin (LTC), bitcoin cash (BCH), bitcoin SV (BSV), cardano (ADA), chainlink (LINK), binance coin (BNB), crypto.com coin (CRO), EOS, DASH, Stellar (XLM) and Monero (XMR). The sample periods start in 2015 at the earliest and often end in 2020 or 2021. The most common types of data examined are crypto asset prices or returns and blockchain transaction flows, often aggregated at hourly or daily frequency.

The research questions and hypotheses of the papers shown in Table 3.1 all revolve around the effect of stablecoins – or often tether in particular – on crypto asset prices, but do not always address the issue of market manipulation directly. Kristoufek (2022) asks whether the amount of stablecoins in circulation reflects crypto asset demand and whether it could be supported that increases in stablecoin supply are the cause of crypto asset price bubbles; and similarly, Kristoufek (2021) also examines the dynamic interaction between stablecoins and major crypto assets. In the above articles, the author carefully avoids the issue of potential market manipulation, by mentioning that any manipulative effect from stablecoin issuance – or lack thereof, hinges upon the validity of fiat currency backing reserves, which is not examined by the author; for instance, Kristoufek (2022) mentions that ‘if the backing is valid and existent (but not necessarily a full backing), then the stablecoin influx signals an increased demand in investment in crypto assets’.

Ante et al. (2021a) focus instead on specific large stablecoin transactions on the relevant blockchains, and examine their effect on bitcoin returns and traded volume; they develop multiple hypotheses as to the potential effect of such stablecoin on-chain flows depending on their characteristics: for instance, they propose that abnormal bitcoin trading volume should be lower for stablecoin transfers with high

information asymmetry and higher for transfers with low information asymmetry; and in a similar logic, positive subsequent (negative prior) abnormal bitcoin returns are associated with stablecoin transfers where cryptocurrency exchanges are receivers (senders) and/or stablecoin treasuries are senders (receivers). In a similar framework and methodology, Ante et al. (2021b) examine the effects of stablecoin issuance on crypto asset returns and trading volume. Again, the authors do not address the issue of potential market manipulation directly, and instead develop hypotheses as to the general informational content of stablecoin on-chain transactions.

The article of Griffin and Shams (2020), originally published on SSRN in 2018 is one of the first papers to analyse tether and bitcoin transactions on their respective blockchains and their influence on crypto asset prices in the 2017 bubble; specifically, Griffin and Shams (2020) associate their framework to that of international capital flow studies and ask whether tether supply is ‘pulled’, i.e. driven by legitimate investor demand or ‘pushed’, i.e. used to create inflation in crypto asset prices. Contrary to the articles mentioned previously, Griffin and Shams (2020) are much more direct in framing their research questions and hypotheses along the issue of potential market manipulation from stablecoins – tether in particular. Similar to Griffin and Shams (2020), the research question of Wei (2018) addresses crypto asset market manipulation directly, by asking whether tether issuance impacts bitcoin’s valuation, particularly, during the late 2017 bull market period. In contrast, the hypotheses of Lyons and Viswanath-Natraj (2020) focus primarily on tether’s price peg, and only examine aggregate tether flows to the secondary market tangentially, under a hypothesis that stablecoins are passive with respect to intervention in crypto asset markets. Similarly, Grobys and Huynh (2021) do not focus directly on crypto asset market manipulation, focussing instead on whether jumps in USDT/USD returns are causal for bitcoin returns, in the context of market efficiency.

Having discussed the research questions and hypotheses in the relevant literature, there now follows an examination of methodologies used and potential methodological issues. The quantitative methodologies used to examine the above research questions

and hypotheses are mostly related to the vector auto-regression (VAR) model and/or, in the presence of cointegrating relationships, the vector error-correction model (VECM), while another group of papers employ event study methodologies. The choice of methodology very much depends on the nature of the research questions and hypotheses under examination; however, an event study methodology appears more appropriate, if the purpose is to determine whether stablecoins are the cause of crypto asset market manipulation. Kristoufek (2022) uses the VAR/VECM framework and further focus on analysing the interactions and causality between the market cap of crypto assets and stablecoins, while Kristoufek (2021) follows a similar methodology but instead focuses on the volatility decomposition and spillovers. Wei (2018) also uses the VAR model framework to apply Granger-causality tests on bitcoin returns and tether supply issuance; Grobys and Huynh (2021) also employ VAR with lagged and interaction terms to test for the existence of a causal relationship between tether price jumps and bitcoin prices, where jumps in returns are calculated based on asymptotic distribution theory.⁴ Similarly, Lyons and Viswanath-Natraj (2020) also apply a VAR model on stablecoin supply and crypto asset prices and use the local projections method to calculate the impulse response functions.

Variants of event study methodologies combined with linear regression models are used by Ante et al. (2021*a*), Ante et al. (2021*b*) and Griffin and Shams (2020), all based on blockchain transactions of tether and/or other stablecoins and their effect on the returns of bitcoin and/or other crypto assets; these methodologies and the relevant findings are examined in more detail later in this section. It should also be noted that the analysis on blockchain transactions performed by Ante et al. (2021*a*), Ante et al. (2021*b*), Griffin and Shams (2020), Lyons and Viswanath-Natraj (2020) and Wei (2018) is based on methodologies developed previously in a number of papers that examine blockchain transactions of various crypto assets. For instance Somin et al. (2018) analyse the transaction patterns of ERC-20 tokens on the Ethereum

⁴Grobys and Huynh (2021) calculate jumps as the difference between realized variance and realized bi-power variation calculated using hourly returns, then encoded as binary jump variables based on jump size significance.

platform. Bovet et al. (2018) use blockchain transaction network metrics to identify price bubbles in bitcoin. Maesa et al. (2017) also use network metrics such as in- and out- degree to examine bitcoin user behaviour. Tasca et al. (2018) and Meiklejohn et al. (2013) use network centrality measures and other methods to group bitcoin addresses into user groups and identify the entities and business categories in the bitcoin ecosystem such as miners, gambling services, black markets and exchanges.

With regard to potential methodological issues, in the papers presented in Table 3.1, it is important to make a distinction between the use of daily frequency market data, as employed by Wei (2018), Lyons and Viswanath-Natraj (2020), Kristoufek (2021) and Kristoufek (2022), and hourly frequency data as used by Ante et al. (2021*b*), Ante et al. (2021*a*), Griffin and Shams (2020) and Grobys and Huynh (2021). While daily frequency data may be suitable for use in some cases, such as in volatility analysis applications, it is probably better to employ a higher frequency – such as hourly, when it comes to the detection of potentially distortive stablecoin effects on crypto asset prices. This is in part due to the overall significantly faster pace of trading in crypto asset markets compared with traditional financial markets; for instance, in the context of bitcoin spot and derivatives trading, Alexander and Heck (2020) find that arbitrage opportunities exist for no more than 10-15 minutes, when arising from price deviations between different exchanges due to a 10% hypothetical price shock. Moreover, the majority of stablecoin transactions occur on the Ethereum and TRON networks/smart contract platforms, for which the average block creation time varies between [12-14 seconds](#) for Ethereum and [3 seconds](#) for TRON; even when accounting for the confirmation lag commonly employed by crypto asset exchanges, i.e. the creation of an additional 20-35 blocks on top of the particular block containing the transaction of interest, this yields a transaction execution/finalization time between 1-10 minutes.

Another important methodological aspect is the source of data. In most cases the use of traded or averaged prices is suitable. For instance, Lyons and Viswanath-Natraj (2020) and Griffin and Shams (2020) calculate average crypto asset prices based

on traded prices from exchanges such as Bitfinex, Bittrex and Kraken.⁵ Similarly, Grobys and Huynh (2021) obtain USDT/USD and BTC/USD price data from the Bitfinex exchange, and Ante et al. (2021a) and Ante et al. (2021b) obtain price and traded volume data on BTC/USD and other pairs from Bitstamp. However, the use of [Coinmarketcap](#) as a data source requires careful treatment; Wei (2018) makes use of aggregate traded volume data from Coinmarketcap and Kristoufek (2021) and Kristoufek (2022) use the same source to obtain crypto asset and stablecoin market cap data.

While the use of aggregated volume and simple- or volume-weighted average price data is not necessarily problematic, the calculation methodology of such averaged data may cause issues, and Coinmarketcap is a notable case in point, as also highlighted in Alexander and Dakos (2020). The [methodology](#) of Coinmarketcap makes use of cross-rates to infer traded volumes that are then also used in VWAP calculations. For instance, in the calculation of aggregate traded volume and VWAP price for the BTC/USD pair, Coinmarketcap not only includes all traded BTC/USD pairs but also obtains the price and volume of all cross-rates e.g. the BTC/ETH pair and in this case uses the ETH/USD price to convert them to USD, thus assuming that all trades in the BTC/ETH pair are ultimately converted to USD, which is by no means the case. The issue is even more evident when examining traded volume, as a significant overlap can exist in the traded pairs used to calculate aggregate volume for different crypto assets; for instance, the same top three traded pairs are used (and therefore double-counted) for the daily aggregate traded volume (and VWAP price) for [bitcoin](#) and [tether](#) on Coinmarketcap, namely the BTC/USDT spot pair on Binance, OKEEx and Huobi.

The overall findings of the literature shown in Table 3.1 are mixed, with some papers finding that stablecoin flows have a significant effect on crypto asset prices, while several others find no evidence of such a relationship. This observed lack of consensus could stem from differences in the sample periods used, but also from the

⁵Lyons and Viswanath-Natraj (2020) and Griffin and Shams (2020) obtain data from the [CoinAPI](#) provider.

nature of the stablecoin-related data examined. Depending on the methodology used, it is often the case that initial datasets contain a very large number of stablecoin transactions, which the authors then filter based on criteria such as the transaction sender and/or receiver and the transacted amount; these selection criteria could therefore result in the inclusion both of transactions with a manipulative effect and also of transactions that have no significant effect on crypto asset prices. This could in turn produce confounding effects in the methodologies used, resulting in the observed lack of consensus.

Kristoufek (2022) determines that stablecoins issuance is not an ‘ignition point’ of crypto asset market bull runs, and therefore concludes that increasing stablecoin supply mostly reflects an increasing demand for investing in crypto assets. Kristoufek (2021) identifies volatility spillover dynamics between stablecoins and crypto assets, but does not interpret them as signs of a ‘bubble-boosting mechanism’, as the volatility spillovers towards stablecoins drop to zero during the 2019 period of low crypto asset prices, which is contrary to the expectation under the bubble-boosting hypothesis that spillovers both to and from stablecoins should increase. Kristoufek (2021) instead concludes that increased stablecoins supply is a reflection of increasing demand for investing in crypto assets.

Grobys and Huynh (2021) find that positive jumps in USDT/USD returns are Granger-causal for BTC/USD returns, interpreting this as increased demand for USDT causing the USDT/USD jump, then followed by large sell volume in the BTC/USDT pair and resulting in lagged price drops; they also find that positive USDT/USD return jumps are associated with contemporaneous large negative BTC/USD returns. Wei (2018) finds no evidence that tether issuance causes significantly positive subsequent bitcoin returns, while traded volume is also unable to predict subsequent bitcoin returns; however, tether issuance events are followed by subsequent increased traded volume for both bitcoin and tether, and tether volume increases following days with negative bitcoin returns; the author also finds strong autocorrelation in the time series of daily tether issuance, indicating that large tether

issuance volume is spread out across several days either intentionally or due to time clustering in the demand for tether. Lyons and Viswanath-Natraj (2020) find that the impulse response functions of bitcoin and ether prices are not statistically different from zero at a 5% significance level and at a horizon of up to 60 days; put simply, the authors find that bitcoin and ether prices do not respond to tether flows from the tether treasury towards the secondary market, and therefore conclude that there is no evidence of stablecoin issuance driving crypto asset prices, instead finding that stablecoin issuance responds to deviations of their dollar prices from the pegged rate.

In contrast with the above articles which find no significant effect of stablecoins on crypto asset prices, the findings of Ante et al. (2021*a*), Ante et al. (2021*b*) and Griffin and Shams (2020) are partially or fully in support of the hypothesis that stablecoin flows and/or issuance have a significant effect on crypto asset prices. Ante et al. (2021*a*) find significant abnormal BTC/USD returns over the 12 hours before a stablecoin transfer. When broken down by transaction type, abnormal returns are observed before transactions that originate from stablecoin treasuries, but not after the transactions occur, except when flows are between unknown addresses. Negative price effects occur prior to transfers to treasuries, due to the withdrawal of capital from the crypto asset market and similarly, transfers between two crypto asset exchanges are associated with negative returns. Similarly, Ante et al. (2021*b*) find that crypto asset returns are significantly negative in the periods before stablecoin issuances, while they stabilize with the issuance events.

Griffin and Shams (2020) find that only on days after new tether supply is issued, bitcoin-tether concurrent flows from Bitfinex to Bittrex and Poloniex increase just below round number thresholds in the bitcoin price but drop immediately after the threshold; the authors find no such evidence in the bitcoin-tether flows to and from other tether exchanges such as Binance, Huobi or OKEx. They further find that bitcoin-tether flows below the threshold are significant only on days following new tether issuance, indicating attempts to halt the downward price trend and, through investor herding behaviour, create a price support level. A significantly

higher ‘future’ bitcoin return is identified only on periods following tether issuance and negative lagged bitcoin returns and bitcoin end-of-month daily returns are negative and significantly lower in months with high tether issuance. The authors observe significant reductions in Bitfinex’s bitcoin balance at the end of months with high tether issuance, while finding no such effect on the bitcoin balances of the top-20 crypto asset exchanges. The authors also explore alternative explanations and hypotheses that may explain their findings, such as that demand for tether is driven by fiat currency-holding investors seeking to invest in bitcoin, but find no significant evidence to support them.

Having discussed the research questions, methodologies and findings of the relevant papers shown in Table 3.1, there now follows a focussed discussion of the article by Ante et al. (2021*a*), which is the most relevant as to the research question and methodological framework examined in this Chapter, followed by a presentation of the incremental contributions this Chapter introduces.

Ante et al. (2021*a*) use an event study methodology to calculate abnormal returns and abnormal trading volumes potentially caused by blockchain stablecoin transfers, based on cross-sectional studies e.g. by Chae (2005), Brown and Warner (1985) and Armitage (1995). The authors distinguish between stablecoin transactions sent to or from unknown addresses, crypto asset exchanges and stablecoin treasuries, and suggest that they have varying effects on crypto asset prices. However, they do not account for ‘chain-swap’ transactions, i.e. transactions that convert large stablecoin quantities from one blockchain to another, which are quite common e.g. for tether (USDT) which exists on multiple smart contract/token platforms; this may be reducing the significance of results, as such transactions will register as events in the methodology even though their purpose is purely technical, e.g. in cases where one blockchain is highly congested and users wish to avoid high transaction fees, so chain-swaps are not expected to have any price impact.

The expected return and trading volume are calculated by Ante et al. (2021*a*) as the average over an estimation window that covers an observation period of

150 hours to 15 hours before each stablecoin transfer, i.e. they are employing the constant mean model. The expected return and volume are then compared with the corresponding observed return and volume during a 25-hour event window, from -12 hours to +12 hours around each stablecoin transfer event, while different windows are also examined for robustness. The abnormal return and volume are then calculated as the difference between the expected and the observed return and volume, and their significance is evaluated using t-tests and also the non-parametric Wilcoxon signed rank test which assumes an equal likelihood of positive and negative abnormal returns and volumes.

Having described the methodology of Ante et al. (2021*a*) and also of Ante et al. (2021*b*) who follow a similar methodology, the discussion now focuses on the event and event window design, a crucial methodological point, given the distribution of stablecoin blockchain transaction flows. Event studies in traditional finance often examine clearly-defined and discrete events such as earnings announcements or stock splits; however, when the events examined are stablecoin blockchain transaction flows which occur very frequently and are often clustered across time, a fixed-length event window may sometimes either include unrelated returns in the event period or omit clustered events that fall just outside the event window; similarly, overlaps in both the event windows and observation periods may exist, even when aggregating the data to an hourly frequency, so that e.g. the estimated effects of one event may occur in the observation window of another event. These types of ‘contamination’ are to be avoided if possible, to ensure that the abnormal effects can be fully explained by the event examined; for instance Hakala (2017) includes additional controls for the effect of past events in the estimation period. These issues are also acknowledged by Ante et al. (2021*a*), who argue that such issues cannot be avoided, as their dataset contains no non-overlapping events.

Another important issue in the methodology of Ante et al. (2021*a*) is the use of the constant-mean model. While the constant-mean assumption may be suitable for traded volume, it does not have sufficient explanatory power for crypto asset

returns and therefore its ability to extract the abnormal return during an event may be limited. For instance, a simple experiment using hourly BTC/USD historical price data can show that the explanatory power of the constant-mean returns model even on 1-hour ahead BTC/USD returns is usually quite small. Obtaining hourly BTC/USD log returns based on historical price data from Bitstamp between April 2019 - March 2020, as used by Ante et al. (2021a), it is easy to calculate the 150-hour moving average of the returns. The above sample period can then yield 8,636 estimations of the constant-mean returns model, i.e. a simple linear regression with BTC/USD hourly log returns as the dependent variable and an independent variable calculated as the 1-hour lag of the 150-hour moving average on BTC/USD log returns, using 149 sample points in the estimation of each regression model. Given the above, the estimated regression R^2 is at most 0.22 but in 99% of estimations it is even below 0.08, with a fairly low average of 0.014 and a median of 0.008.

Finally, there follows an overview of the incremental contributions this Chapter introduces to the crypto asset financial literature. Firstly, this is the first empirical study to examine the materiality of tether token flows originating from the tether treasury in causing positive abnormal bitcoin returns and ultimately estimating damages incurred by the (allegedly) manipulative practices. The contribution of this Chapter is denoted by the fact that a number of courts have rejected or refused to admit into evidence damages reports or testimony by damages experts in securities cases which fail to include event studies or ‘something similar’ (Fisch et al., 2017).

The methodology presented in this Chapter is adapted from the standard framework of securities fraud litigation event studies, and the hypothesis of significant positive abnormal bitcoin returns is examined with both parametric and non-parametric tests for robustness, while the materiality of each USDT outflow event from the tether treasury towards crypto asset exchanges and other unknown entities is examined separately, to account for the possibility that a small number of potentially manipulative events are ‘concealed’ amongst a much greater number of unrelated events.

The effect of past events in the estimation period is controlled for, as suggested in the relevant literature, and abnormal returns are estimated using index model specifications instead of the constant-mean returns approach. A flexible event window size is employed to account for clustering in USDT outflow events from the tether treasury, and chain swap transactions – that are executed in order to swap USDT tokens between token platforms for rebalancing purposes and/or to avoid network congestion – are excluded from the sample, as they are not related to market movements.

The retrieval of transaction data from the Bitcoin/Omni Layer blockchain is based on our own experience of establishing a local blockchain node and is examined in significant detail, which also constitutes an important contribution, as most researchers only use secondary blockchain data sources for crypto asset financial research and any academic papers that make use of data obtained directly from a blockchain do not provide detailed explanations on how this is achieved.

3.3 Methodology

This Section provides an overview of the econometric methodologies used to determine the effect of USDT outflows from the tether treasury towards crypto asset exchanges and other unknown entities on the price of bitcoin. As the methodology used is adapted from securities fraud litigation event studies, there follows a brief discussion which motivates this choice; subsequently an overview of the methodology is provided, together with a clarifying differentiation from the standard cross-sectional event study framework, followed by the event window construction, the index model configuration and estimation, the abnormal returns hypothesis test and the price effect and damages estimation.

As to the motivation of adapting the securities fraud litigation event study framework, Dove et al. (2019) highlight the requirement by courts that plaintiffs in a securities case should demonstrate the statistical significance of an event-date

return, so this methodology is directly related the materiality inquiry in securities fraud litigation, for which Baker (2016) mentions that it ‘should involve a fact-specific analysis of whether a reasonable investor would hold the particular alleged misrepresentation or omission to be significant in the context of the information available to the market’. After the abnormal return and price effect are determined, Baker (2016) also suggests that damages can then be estimated on an out-of-pocket basis, i.e. ‘each purchaser of a security is entitled to the difference between the price paid for the security and the price it would have traded at, had there been no fraudulent misrepresentation or omission’.

The overall methodology employed in this Chapter consists of adapting the single-company event study framework of securities fraud litigation studies, as described e.g. by Gelbach et al. (2013), where an index model is used to determine the expected return based on a broad market index and estimate the abnormal return caused by the event(s) examined. The hourly return of the BTC/USD price is the dependent variable in the regression model, and the expected return is determined based on an index model estimated via linear regression, where various factors are considered such as the hourly return on an equally-weighted market portfolio based on large population of liquid crypto assets, the return on the cap-weighted MVIS Digital Assets (MVDA) index, and also principal components extracted from the correlated system of crypto asset returns.

The events examined are binary variables based on the occurrence of USDT outflows from the tether treasury towards crypto asset exchanges and other unknown entities, aggregated at an hourly frequency. As USDT outflows can be related to different market players, e.g. as described in the investigative article of Protos (2021a), events are examined on an individual basis and materiality is determined separately for each one. Once material events are determined, the price effect is estimated based on the average abnormal return in each case, and is then combined with bitcoin’s traded volume to produce an estimate of the damage caused by the USDT flow in question, as described by Baker (2016).

It is important to note that while this methodology is somewhat similar to the ‘skeletal econometric structure of an event study’ described by Corrado (2011), it is not directly relevant to the well-established cross-sectional event study methodology as developed by Brown and Warner (1985) and examined by Armitage (1995). The methodology presented in this Chapter differs from the standard cross-sectional event study in several aspects. For instance, a single asset (bitcoin) is examined, so that the significance of abnormal returns cannot be assessed using test statistics that correct for cross-sectional variance. Also, the events examined (USDT outflows from the tether treasury) are highly clustered and with non-standard duration, so that the common fixed-length event window approach employed e.g. by Ante et al. (2021a) is less preferable compared with a flexible variable-length window approach. For the same reason, the model estimation period very often contains previous events, which Ante et al. (2021a) acknowledge as unavoidable in their methodology, while the methodology presented in this Section can account for this feature with the inclusion of an additional control variable in the model.

As the purpose of this study is to determine whether and which USDT outflows from the tether treasury cause positive abnormal bitcoin returns, the USDT outflow events are examined separately, aggregated at the hourly frequency. Contrary to the standard practice of defining a fixed-length event window around the time of the event, a more bespoke approach is chosen to mitigate certain issues that occur in the fixed-length event window approach, stemming from the fact that several USDT outflow events span multiple consecutive hours and many such events are clustered with as little as a single hour in-between. The approach chosen treats USDT outflows from the tether treasury that occur on consecutive hours as a single event. Additionally, and as it is reasonable to assume that the price effects of a USDT outflow may not manifest within the same hour of the event and instead begin slightly earlier e.g. due to insider information or with a slight delay e.g. due to retail investor herding, it is important to include 1 hour before and after each USDT

outflow to the corresponding event window; for this reason, USDT outflow events that are 1 or 2 hours apart are also treated as a single event.

Following the overview of the methodology, we now formalize the index model specification used to estimate abnormal returns. For each event i starting at time t_i^{start} and ending at t_i^{end} and for all t within the m -hour estimation period of event i , i.e. $t \in [t_i^{end} - m, t_i^{end}]$ – where $t_i^{end} - m < t_i^{start} < t_i^{end}$, the return r_{it}^{BTC} on the BTC/USD average (VWAP) price following the specification of Gelbach et al. (2013) is modelled as:⁶

$$r_{it}^{BTC} = \alpha_i + \beta_i X_{it} + \gamma_i D_{it} + \delta_i D_{it}^{(past)} + \varepsilon_{it}, \quad (3.1)$$

where separate regression models based on (3.1) are estimated for each event i to account for the possibility that not all of the events may be material, and for the regression model relating to each event i , the sample period is $t_i^{end} - m \leq t \leq t_i^{end}$; X_{it} represents the market factor discussed below and β_i the corresponding sensitivity; D_{it} is the binary variable set to 1 for all hours included in event i , i.e. for $t_i^{start} \leq t \leq t_i^{end}$ and γ_i is the average abnormal return over the event period of event i ; similarly $D_{it}^{(past)}$ is equal to 1 for all previous events $j = 1, \dots, i - 1$ included the estimation period $t \in [t_i^{end} - m, t_i^{end}]$ of event i and δ is the corresponding average abnormal return, so the model specification controls for the effect of past events in the estimation period as suggested by Hakala (2017), reducing the impact of previous unrelated events occurring during the estimation period.

Regarding the market factor X_{it} in model (3.1), ideally a multi-factor model extension of (3.1) would be used, such as the Fama-French 5-factor or the Carhart 4-factor model. However, such well-developed factors do not exist for crypto assets, and the corresponding equity-based factors are unfit for use as they refer to a completely different market. The main choice for the market factor is therefore based on the return of an equally-weighted ‘market portfolio’ consisting of all crypto assets with

⁶Note that the inequality $t_i^{end} - m < t_i^{start} < t_i^{end}$ holds due to the choice of m to be 720 hours, while none of the events in the sample last for more than 6 hours.

available data and sufficient liquidity. Additional choices are also examined, such as a market factor proxied by the return on the cap-weighted MVDA index which contains the top 100 crypto assets, and also multiple factors in a multi-factor extension of (3.1), obtained as the first few principal components from the correlated system of all crypto asset returns with available data and sufficient liquidity.

The model specification in (3.1) can be estimated using OLS linear regression, regardless of the market factor variant chosen. The estimation period for the regression model of each event i ending at time t_i^{end} is chosen to include all hourly data in the 30-day period prior to the event's end, i.e. $m = 720$ hours, so the estimation period for each event i includes all $t \in [t_i^{end} - m, t_i^{end}]$. This choice of m is made so that the estimation period is large enough to allow for the proper estimation of the model specification in (3.1) but small enough to avoid including confounding effects from previous extreme returns.

The materiality of each USDT outflow event i in causing positive abnormal bitcoin returns is determined based on the estimated average abnormal return $\hat{\gamma}_i$ defined in equation (3.1). A null hypothesis of a zero abnormal return is examined versus the one-sided alternative that the abnormal return is positive. The one-sided hypothesis test is better suited to our research question of whether USDT outflows from the tether treasury are the cause of positive abnormal returns for bitcoin. Moreover, Dove et al. (2019) also argue that 'a one-tailed test is unquestionably more appropriate in securities litigation as it improves statistical power and reflects that the null hypothesis being tested is usually one-tailed'.

The simplest way to perform this hypothesis test is assuming that estimated residuals of model (3.1) follow an independent and identically distributed (i.i.d.) process and are also normally distributed. In this case, the hypothesis test for the significance of the estimated average abnormal return $\hat{\gamma}_i$ for event i is:

$$H_0 : \frac{\hat{\gamma}_i}{est.s.e.(\hat{\gamma}_i)} = 0 \text{ vs. } H_1 : \frac{\hat{\gamma}_i}{est.s.e.(\hat{\gamma}_i)} > 0, \quad (3.2)$$

where the $\frac{\hat{\gamma}_i}{est.s.e.(\hat{\gamma}_i)}$ ratio is compared against the $\alpha\%$ left-tail critical value of the standard normal distribution.

However, the normal i.i.d. assumption for the residuals may be considered weak, so the non-parametric SQ (sample quantile) test is also examined, as described by Fisch et al. (2017) and Dove et al. (2019). The SQ test does not involve any distributional assumptions for the residuals; in the raw residuals version of the SQ test, the estimated average abnormal return $\hat{\gamma}_i$ for each event i is compared against the $\alpha\%$ right-tail quantile of the estimated raw residuals. Similarly, the standardized version of the SQ test is performed similar to equation (3.2) using the standardized estimated abnormal return $\frac{\hat{\gamma}_i}{est.s.e.(\hat{\gamma}_i)}$ compared against the $\alpha\%$ right-tail quantile of the standardized estimated residuals, i.e. the residuals scaled by the regression standard error, per the standard practice as mentioned e.g. by Fisch et al. (2017).

Additionally, given the BTC/USD price $P_{t_i^{start}}$ at time t_i^{start} when USDT outflow event i begins, an estimate of the price effect for the event can be obtained using the corresponding estimated average abnormal return $\hat{\gamma}_i$ from (3.1) as:


$$P_i^{Effect} = \hat{\gamma}_i P_{t_i^{start}} . \quad (3.3)$$

As mentioned previously, the price effect of each event can in turn be combined with the corresponding bitcoin traded volume to produce an estimate of the damage caused by the USDT outflow event i as:

$$Damage_i = \frac{V_i^{BTC}}{2} P_i^{Effect} , \quad (3.4)$$

where V_i^{BTC} is calculated as the total traded volume – expressed in BTC, of all BTC/USD and BTC/USDT traded spot pairs at t_i^{start} , and divided by 2 to account for the fact that only short positions in the corresponding traded pairs would be negatively affected from a positive abnormal bitcoin return which causes a stabilization and even a reversal of a downwards price trend.

Figure 3.1: Hourly BTC/USD price and USDT supply



Hourly BTC/USD VWAP price obtained from Cryptocompare (upper panel) and USDT supply constructed based on blockchain transaction data (lower panel). The sample period is 4 April 2019 - 1 January 2021.

3.4 Data

This Section provides an overview of the data used for the application of the methodology shown previously, and also a detailed discussion on data acquisition. Overall, the sample period examined at the hourly frequency is 4 April 2019 - 1 January 2021, with the sample period's start determined by the availability of relevant blockchain transaction data. The sample contains two periods of significant increases in bitcoin's price, together with an increase in USDT supply from approximately 2 billion tokens to 20 billion, as shown in Figure 3.1.

As both market data and blockchain transaction data are required, each type is discussed in separate subsections. Additionally, the acquisition process of blockchain data is described in great detail, as they are an uncommon data type; a separate analysis of the installation and data acquisition process is presented for the Omni Layer blockchain node, the first platform – second-layer protocol on top of the Bitcoin blockchain, on which USDT tokens were deployed. This constitutes an important contribution as this method is employed by very few academic sources in the crypto

asset finance literature and never described in detail, with most relevant papers resorting to online blockchain explorers.

3.4.1 Market Data

Hourly frequency price and traded volume historical data are obtained from Cryptocompare for the sample period of 3 March 2019 - 1 January 2021, i.e. with the period's start set one month prior to the earliest available blockchain transaction to account for the 1-month estimation period in the models described in Section 3.3. Cryptocompare is selected as the primary market data source, due to its reliable price calculation methodology, as explained in Alexander and Dakos (2020). The primary crypto asset examined is bitcoin, and additional hourly price data are obtained on the cap-weighted MVIS Digital Asset (MVDA) index containing 100 crypto assets and also on a large number of crypto assets to construct the equally-weighted market portfolio. Specifically, the return on the equally-weighted market portfolio is constructed as the average return on the USD-denominated prices of all crypto assets ranked by market cap that exhibit a traded 24-hour volume of more than \$1 million on 1 January 2021 and have available historical data going back to the sample period's start. Additionally, all stablecoins and other price-pegged crypto assets are excluded as well as any cases that exhibit extreme outliers, yielding a total of 135 crypto assets included in the market portfolio.⁷

Table 3.4 presents the sample statistics for the ordinary returns on: bitcoin's price, the equally-weighted market portfolio based on 135 crypto assets, the MVDA index

⁷The crypto assets included in the market portfolio are: BTC, ETH, XRP, BNB, ADA, LINK, DOGE, XLM, LTC, BCH, TRX, VET, THETA, EOS, NEO, BSV, MIOTA, XMR, OKB, XTZ, SNX, ETC, XEM, NEXO, MKR, DASH, ENJ, ZIL, ZEC, MANA, DCR, KCS, ICX, BAT, ONT, ZRX, RVN, QTUM, WAVES, IOST, SC, BNT, DGB, ZEN, OMG, DENT, FTM, CHSB, UQC, STORJ, LRC, NANO, KNC, VTHO, STMX, BCD, NKN, POWR, AGI, FUN, IOTX, ELF, STEEM, ARDR, POLY, KMD, WAN, ETN, BTM, REP, PPT, UTK, DNT, STRAX, MTL, ELA, SYS, BLZ, AST, QKC, GAS, RLC, VITE, RDD, AION, ADX, DATA, LOOM, REQ, NRG, DBC, NULS, FIRO, RCN, MITH, PIVX, RDNN, GRS, DOCK, QSP, GTO, NAS, CND, YOYOW, SKY, PRO, NXT, NEBL, AMB, NCASH, QLC, GO, SWFTC, GVT, WABI, CDT, WTC, WPR, EGT, SNGLS, TRUE., VIA, SNM, OST, EVX, MDA, POA, DLT, PAY, VIB, MTH, PST, STK, CTXC and UBEX. The stablecoins and other price-pegged crypto assets excluded from the index are: Tether, USD Coin, Wrapped Bitcoin, BUSD, Multi Collateral Dai, TerraUSD, Paxos Standard, True USD, USDJ, sUSD, Tether Gold, Gemini Dollar, tBTC and Hive Dollar.

and also the first four principal components extracted from the correlated system of the 135 crypto assets included in the market portfolio. The equally-weighted market portfolio and MVDA index returns both exhibit similar behaviour compared with bitcoin returns, with slight differences in the extreme values, as shown by the skewness and excess kurtosis statistics; this can also be observed in the corresponding returns time series shown in Figure 3.2. The principal components shown in Table 3.4 are calculated based on the entire sample period, explain approximately 30% of total variance in the correlated system of returns and are shown here for illustrative purposes, as the principal components are calculated separately for each USDT outflow event based on the crypto asset returns data from the corresponding 30-day estimation period. Compared to bitcoin and the two market indices, the principal components exhibit significantly different behaviour, as also shown in the corresponding time series included in Figure B1 in Appendix B, indicating that their use is perhaps less advisable.


Table 3.2: Summary statistics of hourly returns

	BTC	Eq. w. Portfolio	MVDA	PC1	PC2	PC3	PC4
Mean (%)	0.02%	0.05%	0.01%	0.73%	-0.75%	-0.05%	-0.24%
St. Dev. (% p.a.)	72.60%	77.87%	70.12%	1255%	1015%	993%	884%
Skewness	0.43	-1.51	-0.68	7.34	-3.29	2.45	1.18
Ex. Kurtosis	76.5	35.1	53.8	114	37.6	51.6	26.3
Min.	-15%	-16%	-15%	-75%	-169%	-142%	-150%
Max.	20%	14%	15%	313%	79%	181%	161%

Summary statistics of hourly arithmetic returns on bitcoin (BTC), the equally-weighted market portfolio based on 135 crypto assets, the cap-weighted MVDA index and also the first four principal components extracted from the correlated system of the 135 crypto assets included in the market portfolio, based on VWAP USD-denominated prices obtained from Cryptocompare. The sample period is 3 March 2019 - 1 January 2021. The mean, minimum and maximum are expressed in % and the daily standard deviation is annualized using a factor of $\sqrt{24 \times 365}$.

Similarly, Table 3.3 presents a correlation analysis for bitcoin, the equally-weighted market portfolio and the cap-weighted MVDA index, with the principal components again included for illustrative purposes. It is observed that the returns on the equally-weighted market portfolio are highly correlated (0.78) with bitcoin returns, suggesting that it is suitable for use as a market index in model (3.1) described

Figure 3.2: Hourly returns



Hourly arithmetic returns on bitcoin, the equally-weighted market portfolio based on 135 crypto assets and the cap-weighted MVDA index, based on VWAP USD-denominated prices obtained from Cryptocompare. The sample period is 3 March 2019 - 1 January 2021.

previously. The same observation can be made for the MVDA index, however, given that it is calculated using a cap-weighted methodology which produces significant weighting bias for a very small number of large cap crypto assets, including bitcoin, the equally-weighted approach is considered as preferable. Again, the principal components exhibit significantly different behaviour compared with the other data and are, by construction, uncorrelated with each other.

Table 3.3: Correlation matrix of hourly returns

	BTC	Eq. w. Portfolio	MVDA	PC1	PC2	PC3	PC4
BTC	1	0.78	0.94	0.07	-0.29	0.04	-0.71
Eq. w. Portfolio	0.78	1	0.82	0.22	-0.49	0.03	-0.84
MVDA	0.94	0.82	1	0.08	-0.31	0.03	-0.74
PC1	0.07	0.22	0.08	1	0	0	0
PC2	-0.29	-0.49	-0.31	0	1	0	0
PC3	0.04	0.03	0.03	0	0	1	0
PC4	-0.71	-0.84	-0.74	0	0	0	1

Correlation matrix of hourly arithmetic returns on bitcoin (BTC), the equally-weighted market portfolio based on 135 crypto assets, the cap-weighted MVDA index and also the first four principal components extracted from the correlated system of the 135 crypto assets included in the market portfolio, based on VWAP USD-denominated prices obtained from Cryptocompare. The sample period is 3 March 2019 - 1 January 2021.

3.4.2 Blockchain Data

Hourly frequency USDT outflows from the tether treasury are examined between 4 April 2019 - 1 January 2021, based primarily on data retrieved from the Whale Alert online service. Whale Alert is chosen as a one-stop solution that covers multiple crypto assets but only includes very large transactions.⁸. In the following, the blockchain transaction data retrieval problem is described specifically for the case of tether, which is deployed on a multitude of blockchains and smart contract platforms; after motivating the use of a whale alert service as a data source, there follows a description of the data retrieval process and a presentation of the relevant sample statistics.

⁸The term ‘whale’ refers to entities with significant crypto asset holdings that often execute very large transactions

As of the sample period's end on 1 January 2021, tether is deployed as a token on the following blockchain protocols: the Omni Layer, a token platform running on top of the Bitcoin protocol, since 2014; Ethereum, as an ERC-20 token since 2017; TRON as a TRC-20 token, EOS and Liquid since 2019; Algorand and Simple Ledger Protocol (SLP, a token platform on top of the Bitcoin Cash protocol) since 2020. The direct approach for retrieving a full transaction history for tether is either to find an online data provider that can offer this data reliably or set up full nodes for all blockchains on which tether exists.

The tether token (USDT) is deployed on multiple platforms, very possibly to avoid congesting a single token platform and the higher transaction fees that would result from prolonged congestion. As of the sample period's end on 1 January 2021, only the Omni Layer, Ethereum and Tron tether tokens have sufficient supply and transaction volume to be considered important. Given the above, it would be necessary to obtain relevant USDT transactions from the Omni Layer, which is discussed in the next subsection, and also use Ethereum and Tron online blockchain explorers similar to Lyons and Viswanath-Natraj (2020) or establish full nodes for both of these blockchains.

Moreover, USDT supply and transaction volume can increase without significant forward warning on any one of the other four token platforms where a tether token is deployed, or indeed on any other fully decentralized token and smart contract platform such as Waves, NEO, NEM and Stellar and also on crypto asset exchange token platforms such as Binance Launchpad, Huobi Prime, OKEx Jumpstart etc. It is obvious that if significant tether supply and transaction volume appears on any of these token platforms, any effort for tether transaction data retrieval will need to include them as well. It is therefore becoming apparent that such an approach easily scales out of proportion, hence the use of Whale Alert as a data source is warranted.

The colloquial term '[whale](#)' in the crypto asset space refers to entities that hold large proportions of a crypto asset's total supply. Such entities are called 'whales' because they have the potential to cause sudden and massive price moves. Blockchain

transactions from such whales are worth monitoring especially if the funds are sent to an exchange, as this can be an indicator of imminent large price movements. For instance, 8,000 bitcoin were deposited on the BitMEX crypto asset derivatives exchange immediately before a Distributed Denial of Service (DDoS) attack on the same exchange, which coincided with a massive drop in bitcoin’s price on 13 March 2020.⁹ With this rationale, so-called ‘whale alert’ websites/applications have emerged, monitoring multiple blockchains and alerting subscribers of large blockchain crypto asset transactions, often with bot-generated messages on social media platforms and messaging applications such as Twitter, Telegram and Discord.

Two important issues when examining blockchain transaction data are entity identification and false-positives. As crypto asset wallet addresses are pseudonymous and multiple addresses can be owned by the same entity, there is the obvious problem of identifying ‘who-is-who’ on the blockchain; unless entities identify themselves as the owners of particular addresses, researchers must resort to clustering algorithms and heuristics to infer address ownership. False-positives occur when very large blockchain transactions are detected that are not related to imminent market movements but instead serve different purposes such as multi-address wallet rebalancing and, in the case of USDT, chain swap transactions in which USDT tokens are swapped between different token platforms via the tether treasury; importantly, crypto asset exchanges often rebalance their cold-storage wallet addresses, producing such false-positive whale transactions.

There exist several whale alert applications, such as [Whale Alert](#), [Bitcoin Block Bot](#) and [Crypto Whale Bot](#). These applications monitor several crypto assets including the large-cap bitcoin, ether and ripple and also all significant tether tokens on the Omni Layer, Ethereum and Tron networks. Such applications also attempt to identify known entities such as crypto asset exchanges and crypto custodians that perform large blockchain transactions. This approach therefore bypasses the issue of monitoring multiple data sources and also the issue of entity identification.

⁹See Dakos (2020) for a Medium blog post that provides additional detail.

Whale Alert is chosen as the source for blockchain transaction data, since it also attempts to filter out false-positive whale transactions.¹⁰ The Whale Alert website provides a detailed [definition](#) of a ‘large transaction’ separately for each crypto asset monitored, also distinguishing between ‘known’ transactions where at least one of the counterparties are known and ‘unknown’ transactions; for instance USDT transactions are classified as large if they exceed \$10 million in dollar value for known transactions and \$20 million for unknown transactions. Using this data source is therefore analogous to the methodology of Ante et al. (2021a) and Ante et al. (2021b) who only examine stablecoin transactions that exceed \$1 million in value.

Whale Alert provides an [API](#) for data retrieval but only provides limited data for free. It is therefore chosen to retrieve the alert messages sent on the Whale Alert [Telegram channel](#).¹¹ The export channel history feature of the Telegram desktop client is used to save the entire history of the channel in html files; the BeautifulSoup Python library is then used to scrape the html files and save the transaction data in tabular form. Note that the Whale Alert Telegram channel history goes back until 4 April 2019, which is a reasonable date to begin the sample period, given that USDT supply started to increase near that time.¹²

Having discussed the motivation of using a whale alert service as the primary source of blockchain transaction data, and also the data retrieval process, there now follows a presentation of the relevant sample statistics. The data retrieval yields an initial dataset of 16,813 crypto asset transactions reported by Whale Alert on Telegram from 4 April 2019 until 1 January 2021, of which 3,425 are USDT transactions. As the research question examined in this Chapter focuses on the effect of USDT outflows from the tether treasury, the above sample is further limited to 1,196 transactions originating from the tether treasury addresses. Moreover, 150

¹⁰For instance, the other two bot applications, Bitcoin Block Bot and Crypto Whale Bot issued an alert on a massive transaction of 43,995.332 BTC, worth more than \$430 million, between unknown entities on 11 June 2020 at 03:05:16 UTC; Whale Alert recorded the [transaction](#) but did not issue any alert on it because it was identified as an internal rebalancing, with no change in fund ownership before and after the transaction.

¹¹Whale Alert also maintains a [Twitter profile](#) where all transaction alert messages are posted, but retrieving data from Twitter is avoided, as this process has become notoriously difficult.

¹²See e.g. Alexander and Dakos (2019) for a Medium post in 2019 about increasing tether supply.

transactions are identified as chain-swaps, i.e. when USDT tokens are swapped between different token platforms via the tether treasury, and are excluded from the sample. This yields a final sample of 640 USDT transactions sent from the tether treasury to crypto asset exchanges and a further 406 transactions sent from the tether treasury to other unknown entities; the exchanges-unknown entity distinction is followed, similar to the approach of Ante et al. (2021a). The dataset is then re-sampled and aggregated to an hourly frequency, for which sample statistics are shown in Table 3.4.

As shown in Table 3.4, USDT outflows from the tether treasury towards crypto asset exchanges and other unknown entities total approximately 17 billion tokens, with an average transaction size of 19 million towards exchanges and slightly smaller (16 million) towards other unknown entities, and a corresponding maximum of 300 and 125 million respectively. Figure 3.3 shows the distribution of USDT outflows from the tether treasury towards different entities, where more than 90% of outflows are either directed towards Binance or unknown entities, justifying the exchanges-unknown entity distinction used.


Table 3.4: Summary statistics of USDT outflows from the tether treasury

	Exchanges	Unknown
N	547	365
Total	10,778,372,775	5,973,600,014
Min.	1,498,500	2,797,200
Max.	300,000,010	125,000,000
Mean	19,704,521	16,366,027
St. Dev.	22,388,364	15,270,564
Skewness	8.52	3.09
Ex. Kurtosis	97	13

Summary statistics of USDT outflows (expressed in number of USDT tokens, approximately equal to USD) from the tether treasury towards crypto asset exchanges and other unknown entities, aggregated at the hourly frequency. The sample period is 4 April 2019 - 1 January 2021.

As mentioned previously in Section 3.3, USDT outflows from the tether treasury that occur on consecutive hours are each considered as a single event. Out of 491 USDT outflows towards exchanges, 442 occur within a single hour, 42 occur within 2

Figure 3.3: Distribution of USDT outflows from the tether treasury




Distribution of USDT outflows from the tether treasury towards crypto asset exchanges and other unknown entities, during the sample period 4 April 2019 - 1 January 2021. Note that a very small fraction (well below 1%) is also sent to Kraken, Poloniex and Bitbank.

consecutive hours and 8 within 3 hours. Moreover, there are 56 pairs of events that are a single hour apart and similarly 32 outflows are 2 hours apart, which are also each collapsed to a single USDT outflow event, given the addition of 1 hour before and after the USDT outflows to each event window, as also described in Section 3.3. Similarly, out of 325 USDT outflows towards other unknown entities, 293 occur within a single hour, 29 occur within 2 consecutive hours and there is one case each that occurs within 3, 5 and 6 consecutive hours. Moreover, there are 30 pairs of events that are a single hour apart and similarly 25 events are 2 hours apart.


Following the above pre-processing, a total of 402 USDT outflow events from the tether treasury towards exchanges are examined, and similarly 269 USDT outflow events towards other unknown entities. Figures 3.4 and 3.5 indicate the hours on which a USDT outflow event occurs (upper panel) and the size of each outflow in million USDT (lower panel). In both cases, the event clustering is obvious, fully justifying the flexible event window size approach and the inclusion of a control variable for past events in the estimation period, as described in Section 3.3. Additionally, it is observed that both the clustering and size of outflows increase significantly following the Covid-related market crash of March 2020, as also shown previously in Figure 3.1.

Figure 3.4: USDT outflows from the tether treasury to exchanges



USDT outflows from the tether treasury towards crypto asset exchanges. The upper panel shows the flow indicator set to 1 on hours when USDT outflows occur, and the lower panel shows the corresponding flow size in billion USD. The sample period is 4 April 2019 - 1 January 2021.

Figure 3.5: USDT outflows from the tether treasury to unknown entities



USDT outflows from the tether treasury towards other unknown entities. The upper panel shows the flow indicator set to 1 on hours when USDT outflows occur, and the lower panel shows the corresponding flow size in billion USD. The sample period is 4 April 2019 - 1 January 2021.

3.4.3 Tether on the Omni Layer

Having discussed the market and blockchain transaction data, this subsection focuses on the data retrieval process of tether transactions from the Omni Layer, using a full blockchain node that was established locally in a macOS server machine. To the best of our knowledge, only a very small number of the published articles or working papers in the crypto asset finance literature obtain their data directly from a locally installed blockchain node.¹³ A detailed account of this data gathering process is therefore provided, as an aid for further research on this topic.¹⁴ We begin with a description of the data retrieval problem, followed by the approaches followed in the relevant literature and then provide a detailed explanation of the data retrieval methodology developed.

At the beginning of this research, tether only existed on the Omni Layer, a digital token platform deployed as a second-layer protocol on top of the the Bitcoin blockchain.¹⁵ There exists an online [Omni Layer explorer](#) with an application programming interface (API), [OmniAPI](#) and this was considered initially for retrieving tether transactions. However, there are several issues with this data source: it is not possible to retrieve a full transaction dataset, i.e. *all* tether transactions that occurred in a given time frame. Instead, the `OmniAPI /v1/transaction/address` method allows the retrieval of all tether transactions involving a *given* Omni Layer address, which then begs the question of which addresses should be examined. A reasonable approach is to only examine transactions involving crypto asset exchange addresses, some of which can be found in the [tether rich list](#) – a self-reported list on the tether website.

The approach above is followed by Griffin and Shams (2020) in order to obtain tether transactions; they obtained all past versions of this webpage saved on the [Wayback Machine](#), compiled a list of addresses-of-interest and queried OmniAPI for

¹³For instance, Easley et al. (2019) set up a full Bitcoin node to study the evolution of bitcoin transaction fees and Daian et al. (2019) set up multiple full Ethereum nodes to examine frontrunning and other manipulative practices on decentralized crypto asset exchanges deployed on Ethereum.

¹⁴While the primary source of blockchain transaction data used in this Chapter is Whale Alert, the Omni Layer blockchain data retrieval is used for dataset verification purposes.

¹⁵Omni Layer token transactions are embedded inside small value bitcoin transactions.

the transaction history of each of these addresses. While reasonable, this methodology potentially leaves certain gaps, as the information provided by OmniAPI and the tether rich list may not be valid and fully up to date. The OmniAPI data source is probably trustworthy, although it is worth mentioning that Omni Explorer and Omni API are maintained by the Omni Layer development team and some of its members also participate in the founding team of the tether token, indicating a potential conflict of interest. As for the tether rich list, there is again a conflict of interest but more importantly, there is no guarantee that this self-reported list of addresses contains and will continue to contain accurate and complete information and that there are no other important addresses beyond these reported in the rich list.¹⁶ Instead, the approach followed here is to set up a full Omni Layer blockchain node and retrieve the entire tether transaction history from the local copy of the blockchain, as described e.g. by Vicati (2018) in a Hacker Noon blog post.

In the following, we discuss the technical process of setting up a full Omni Layer blockchain node and configuring the client software to execute data retrieval API methods; we then present and compare the available data retrieval methods, we provide a detailed description of the most efficient method's Unix bash script implementation and execution, and finally we describe the structure of the retrieved Omni Layer blockchain transaction data.

We use a macOS server machine and install the Omni Core client that implements an Omni Layer node, using the instructions in the relevant Github [code repository](#). As per the instructions we start up the node using the `omnicored` executable and let it download an up-to-date copy of the Bitcoin blockchain. Tether is an Omni Layer token identified by the unique Omni Layer property ID #31, so we need to retrieve all Omni Layer transactions for a given time frame and filter them to obtain the tether transactions.¹⁷ In order to retrieve Omni Layer transactions, Omni Core includes an API tool, `omnicore-cli` with several methods called RPC

¹⁶Griffin and Shams (2020) overcome this caveat by applying clustering algorithms and identifying other addresses potentially owned by exchanges based on the list of known addresses.

¹⁷The majority of Omni Layer transactions are tether transactions so this is a relatively simple task.

(remote procedure call) commands to access the locally-stored blockchain database.¹⁸ Again following the Omni Core installation instructions on Github, we edit the `bitcoin.conf` configuration file located inside the `/bitcoin` folder created by the installation to include the line `server=1`. We then restart the node, again using the `omnicore` executable file and at this stage we are ready to execute `omnicore-cli` methods.¹⁹

There are several ways to set up a data retrieval routine but as `omnicore-cli` is a command line tool, we choose to retrieve the transaction data with a Unix shell script, also called a bash script. There is no direct `omnicore-cli` method to retrieve the data for all Omni Layer transactions stored in the Bitcoin blockchain; instead a two-step process can be followed: firstly, two methods are provided to retrieve a list of hexadecimal transaction identifiers, henceforth called transaction IDs; we then need to iterate through the list of transaction IDs and call another method that retrieves the data for a single transaction, given its ID.

For the first step, the available methods for retrieving Omni Layer transaction IDs are `omni_listblocktransactions` and `omni_listblockstransactions`. The former takes a single integer number as input and returns a JSON array with the IDs of the transactions included in the block at blockchain height equal to the integer input; the latter takes two integers as input and returns the IDs of all transactions included in blocks inside the blockchain height range designated by the two integers.²⁰

As there is no obvious indication of which retrieval process is the most efficient, we now present alternative data retrieval processes in algorithms 1 and 2, implement them in Unix bash scripts and compare their performance in terms of execution time.

¹⁸As Omni Core is a fork of Bitcoin Core, the main client for the Bitcoin protocol, its functions are very similar to those of Bitcoin Core. See Antonopoulos (2017) for detailed instructions on operating the Bitcoin Core client.

¹⁹See the Omni Core [documentation](#) for a full list of available methods.

²⁰[JSON](#) is a lightweight data-interchange format similar to csv. It contains sequential and/or nested arrays of key-value pairs.

Algorithm 1: Data retrieval with double loop

Input: initialBlockHeight, finalBlockHeight

```

1 for each block between initialBlockHeight and finalBlockHeight do
2   | call omni_listblocktransactions with current block height as
   |   argument
3   for each transaction ID do
4   |   call omni_gettransaction with current transaction ID as argument
5   |   store transaction data into JSON file
6   end
7 end

```

Algorithm 2: Data retrieval with single loop

Input: initialBlockHeight, finalBlockHeight

```

1 call omni_listblockstransactions with initialBlockHeight and
   finalBlockHeight as arguments
2 for each transaction ID do
3   | call omni_gettransaction with current transaction ID as argument
4   | store transaction data into JSON file
5 end

```

In order to determine which algorithm is more efficient, we implement algorithms 1 and 2 as Unix bash scripts and compare their execution times for different block height ranges as input, starting on block 636,791 which was mined on 29 June 2020.


On average, 144 blocks are mined on the Bitcoin blockchain per day but the number of Omni Layer transactions included in each block can vary significantly; in Q2-Q3 2020 there were between 5,000 - 10,000 Omni Layer transactions per day but for instance in Q3 2019 there were as many as 100,000 transactions per day. In figure 3.6 we plot the execution time for the two algorithm implementations (left-hand panel) and also the difference between the execution times (right-hand panel) against the number of transactions retrieved. We observe that algorithm 1 is consistently

faster and as the number of transactions retrieved grows, the difference in execution time tends to increase.

An examination of the algorithms' structure provides some indications as to the reasons for the outperformance of algorithm 1. The main difference in the structure of the two algorithms is in the first of the two `omnicore-cli` method calls where algorithm 1 calls `omni_listblocktransactions` and algorithm 2 calls `omni_listblockstransactions`, so the reason for the difference in execution time is probably related to the construction of these two `omnicore-cli` methods. We propose that algorithm 1 is faster because we access the blocks in the database sequentially, while algorithm 2 seems to access blocks in a somewhat random order.

A detailed review of the Omni Core client source code written in C++ is beyond the scope of this Chapter, so we provide the following details as an aid for further research on this topic: the code segments that define the two methods are located in the `rpc.cpp` file of the Omni Core Github repository; `omni_listblocktransactions` first retrieves the relevant block from the blockchain database and then retrieves the transaction IDs within, while `omni_listblockstransactions` makes use of a function called `GetOmniTxsInBlockRange` which is defined in the `dbtxlist.cpp` file; in turn, the function `GetOmniTxsInBlockRange` makes use of a `leveldb Slice` structure to retrieve transaction IDs from multiple blocks.

Figure 3.6: Omni Layer data retrieval execution times



Left-hand panel: Execution time in seconds for algorithms 1 (double loop) and 2 (single loop), against the number of transactions retrieved from the blockchain. **Right-hand panel:** Difference in execution time in seconds between algorithms 1 and 2, against the number of transactions retrieved from the blockchain.

Regardless of the exact cause for the difference in execution time, the increasing trend in the difference between execution times shown in the right-hand panel of Figure 3.6 indicates that the double loop implementation of algorithm 1 using the `omni_listblocktransactions` method should be used to retrieve Omni Layer transactions in bulk. We now present and describe in detail the bash script that implements algorithm 1 for the retrieval of Omni Layer transactions:

```
#!/bin/bash

cd /Volumes/LocalDataHD/md451/bin/

### Create omni_output.json file to store the tx data
touch omni_output.json

### Put a [ symbol inside the json file to denote
### the beginning of a JSON array
echo "[" > /Volumes/LocalDataHD/md451/bin/omni_output.json

cd /usr/local/omnicore/bin/

### Obtain start and end block for data retrieval
### from user input
start_block=$1
end_block=$2

### Loop through the blocks based on user inputs
for ((i=start_block;i<=end_block;i++)); do

get_tx_ids=$(./omnicore-cli omni_listblocktransactions ${i})
length=$(echo ${get_tx_ids} | jq '. | length')
```

```

### Loop through the tx id's contained in get_tx_ids
for ((j=0;j<=((length - 1));j++)); do

tx_id=$(echo ${get_tx_ids} | jq ".$j]" | tr -d '"')
echo $(./omnicore-cli omni_gettransaction ${tx_id})

### Put a , after each entry unless it is the very last one
if [[ $i -lt $end_block || $j -lt $(((length - 1))) ]]; then
echo ","
fi

done

done >> /Volumes/LocalDataHD/md451/bin/omni_output.json

### Put a ] symbol inside the json file to denote the end
### of a JSON structure
echo "]" >> /Volumes/LocalDataHD/md451/bin/omni_output.json
echo Execution time: $SECONDS seconds.

```

The script begins with the line `#!/bin/bash` indicating to the operating system to invoke the specified bash shell to execute the commands that follow in the script. We then use the `cd` command to change the current directory to the folder where we wish to store the output file of the script, `omni_output.json` and create the file using the `touch` command. With the exception of the first line, all script lines beginning with the `#` character are comments that are ignored in the execution and serve in making the script easier to understand; we designate comments with `###` for increased legibility.

The `echo` command is used to print text to the terminal, but we also add the `>` operator to print the `[` character to the output file; we do so because a JSON file always begins with this character. Following this, we change the current directory to

`/usr/local/omnicore/bin`, where the `omnicore-cli` binary executable is located. The user must determine the initial and final block for the transaction data retrieval by typing the script file name followed by two integers in ascending order on the command line when executing the script. The identifiers `$1` and `$2` refer to these two integer inputs and we use these identifiers to store the user-provided values to the variables `start_block` and `end_block`.

We then build the outer loop where the loop counter `i` iterates through the blocks between `start_block` and `end_block`; the double parentheses permit the expansion and evaluation of arithmetic expressions. Inside the outer loop we call the `omni_listblocktransactions` method for the current block `i`, which returns a JSON array containing each transaction ID enclosed in double quotation marks; we store the JSON array inside the variable `get_tx_ids`, where the `$(...)` structure executes any command within the parentheses in a subshell and returns the output. We also extract the length of the JSON array and store it in the variable `length`, to be used in the inner loop. Here we use the `echo` command in combination with the `|` operator to connect the standard output of one process (the `echo` command) to the standard input of another, the `jq` command; `jq` is a command-line JSON processor which we install from its [Github repository](#) and we use it to extract the length of the JSON array stored inside the variable `get_tx_ids`.

We now introduce the inner loop where the loop counter `j` runs between 0 and `length - 1`, since JSON array entries are numbered starting from 0. Inside the inner loop we first use the `|` operator and `jq` to extract the j^{th} element of the JSON array stored inside `get_tx_ids` and then strip it of the enclosing double quotation marks using the `tr -d '"'` command, storing the transaction ID inside the `tx_id` variable, ready to be used in an `omnicore-cli` method. We then call the `omni_gettransaction` method with transaction ID as an argument and use `echo` to append the transaction data JSON array returned in the output file `omni_output.json`. Then we use another `echo` to append a comma character to the output file in order to separate different JSON array entries; we nest this `echo` inside

an `if` statement that causes all entries to be separated by commas except for the very last entry, which does not need one and for which the outer loop counter `i` is equal to the `end_block` variable and the inner loop counter `j` is equal to `length - 1`.

Here the `echo` commands ‘know’ that we wish to append the output to a file because we include the `>>` operator outside the outer loop, immediately after the second `done` statement that concludes the outer loop. We choose to include the append statement outside the loop instead of including it after each `echo` command to avoid opening and closing the output file multiple times, thereby making the execution a little more efficient. Finally, we append a `]` character to the output file to signify the end of the JSON structure stored inside it and we print the execution time in seconds to the terminal.

We store the entire script in a file, e.g. named `data_retrieval`, and change the file’s access permissions with the `chmod +x data_retrieval` command to make the file executable. To execute the script we need to navigate to the directory where the script file is stored and type `./` on the command line followed by the file’s name and then the two integers that designate the block range for which we wish to retrieve data; e.g. executing `./data_retrieval 636791 636795` on the command line will retrieve Omni Layer transaction data from 5 blocks.

The output file produced by the bash script implementation of algorithm 1 contains multiple entries that contain all the metadata for each Omni Layer transaction, and have the following structure:

```
{
  "txid": "efc7575e06bf5d8bee50b2fbecc79645348820c69234f1b7c9ab485ff106
0e00",
  "fee": "0.00003115",
  "sendingaddress": "1G47mSr3oANXMafVrR8UC4pzV7FEAzo3r9",
  "referenceaddress": "1KaUmoRvKGHJy7y6tMuSbvM1aZMusezrxQ",
  "ismine": false,
  "version": 0,
```

```

"type_int": 0,
"type": "Simple Send",
"propertyid": 31,
"divisible": true,
"amount": "373.30000000",
"valid": true,
"blockhash": "00000000000000000004ca6b53bc9fe222e7a76f67c8876035340e7f
324655f6",
"blocktime": 1593398485,
"positioninblock": 1615,
"block": 636816,
"confirmations": 2577
}

```

The above metadata structure is similar to that of Bitcoin transactions. The **txid** field contains the hexadecimal transaction ID. The **fee** contains the transaction fee in BTC for the Bitcoin transaction that contains the Omni Layer transaction data. The **sendingaddress** is the Bitcoin address of the transaction sender. The field **referenceaddress** has various uses depending on the transaction type declared in the **type** field and encoded in the **type_int** field; for transactions of type Simple Send, the reference address denotes the transaction recipient's address. The **propertyid** is a unique number assigned to each Omni Layer token; 31 corresponds to tether. The **divisible** field shows whether the specific token is divisible and the **amount** is the transacted amount. The **valid** field shows whether the transaction is valid based on address balances and the **blockhash** is the hexadecimal identifier the block containing the transaction. The time that the block was mined is contained in **blocktime**, encoded as a Unix timestamp; in this case this transaction was mined on Monday, June 29, 2020 2:41:25 AM. The position of the transaction inside the block is shown in **positioninblock** and **block** shows the block's height on the Bitcoin blockchain. Finally, **confirmations** indicates how many blocks had been mined on

top of the block containing the particular transaction, at the time of retrieving the transaction data from the blockchain database.

3.5 Empirical Results

This Section presents and discusses the empirical results for the materiality analysis of USDT outflow events originating from the tether treasury and directed towards crypto asset exchanges and other unknown entities. As discussed previously in Section 3.3, the methodology is adapted from the single-company event study model framework for securities fraud litigation studies, assessing the materiality of each USDT outflow even separately to account for the possibility that some – indeed several – events may not be material, and also accounting for the effect of previous events included in the estimation period. The factor model defined previously in (3.1) is estimated via OLS linear regression using hourly frequency data, with an estimation period of 720 hours (30 days), and is implemented in R using custom-written code.

The discussion focuses on results produced from the equally-weighted market portfolio factor model specification, with is considered more suitable compared with the cap-weighted MVDA index and PCA-based specifications, as discussed previously in sections 3.3 and 3.4, while additional results from the cap-weighted MVDA index and PCA-based model specifications are provided in Appendix B. Tables 3.5 and 3.6 present the events and corresponding estimates for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t -test and the raw and standardized SQ tests, to ensure the strictest possible criteria in determining the materiality of USDT outflow events. The t -test critical value is based on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model; similarly, the raw and standardized non-parametric SQ tests are obtained as described in (3.2) of Section 3.3.

Overall, out of a total of 671 USDT outflow events from the tether treasury towards exchanges and unknown entities, 30 events are found to produce a significant positive abnormal bitcoin return and therefore a material price effect. While this result may appear surprising, it is for this reason that the materiality of each event is examined separately, to avoid confounding effects from unrelated events which may in some cases have biased the findings in the relevant literature, as discussed previously in Section 3.2. Moreover, the total estimated damage from the 30 material USDT outflow events amounts to \$54 million, and given that each material event causes a bitcoin price increase which is subsequently amplified by investor herding and fear-of-missing-out behavioural trading patterns, it is very likely that the 30 identified material events may be enough to manipulate bitcoin prices to a significant extent.

Regarding outflows towards exchanges, 19 USDT outflow events towards exchanges are considered material and included in Table 3.5. The significant events are located primarily in Q2-Q4 of 2020 which is considered reasonable given the increase in USDT outflows from the tether treasury in that period as shown previously in Figure 3.4; the average event duration is 2.4 hours, with most events lasting 2 hours and some cases of 4-hour and 6-hour events. The total USDT sent to exchanges in the above events amounts to 407 million USDT, ranging between 5 million - 60 million. The estimated average abnormal return for the above events ranges between 0.28% - 1.38% with an average value of 0.7%; by construction, all $\hat{\gamma}$ estimates included in Table 3.5 exceed their corresponding SQ critical value, and similar for the t_γ and the SQ_{std} and Student- t critical values. The regression adjusted R^2 ranges between 0.06 - 0.67 with an average value of 0.43. The estimated price effect ranges between \$27 - \$268, with an average of \$90. The $P^{Effect}/\Delta P$ ratio, which compares the price effect against the total change in price during each event, often exceeds 1 in magnitude and is often negative, indicating that the estimated price effect caused by the abnormal return is comparable to the price change during event period and is very often in the opposite direction of a downward price movement during the event period. This

suggests that in several cases, the significant abnormal bitcoin return contributed towards a negative price trend reversal. Finally, the total estimated damage amounts to \$23 million, ranging between \$120k and \$6.7 million.

Table 3.5: Material USDT outflows from the tether treasury to exchanges

Start	End	USDT	$\hat{\gamma}$	SQ	t_{γ}	SQ_{std}	Adj. R^2	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-09-19 21:00:00	2019-09-19 23:00:00	4,995,000	0.008	0.005	3.81	1.30	0.52	10,277.6	3.5	85.3	24.6	385,221
2020-05-07 14:00:00	2020-05-07 16:00:00	10,000,000	0.006	0.006	2.45	1.23	0.57	9,506.1	126.6	60.5	0.5	1,037,029
2020-05-07 18:00:00	2020-05-07 20:00:00	12,000,000	0.007	0.006	2.75	1.25	0.56	9,806.1	-3.2	71.7	-22.3	1,266,315
2020-05-27 12:00:00	2020-05-27 14:00:00	24,000,000	0.010	0.009	2.96	1.64	0.53	9,151.1	4.3	90.3	21.2	1,224,616
2020-07-23 16:00:00	2020-07-23 18:00:00	23,000,000	0.005	0.003	4.38	1.33	0.52	9,631.6	-41.0	50.0	-1.2	457,699
2020-07-24 20:00:00	2020-07-24 22:00:00	30,000,000	0.003	0.003	2.39	1.30	0.49	9,627.4	-78.5	27.0	-0.3	119,379
2020-07-26 09:00:00	2020-07-26 11:00:00	10,100,000	0.007	0.003	6.29	1.31	0.43	9,846.6	132.0	72.3	0.5	747,881
2020-07-27 05:00:00	2020-07-27 07:00:00	12,000,000	0.011	0.003	7.42	1.19	0.28	10,299.7	-129.0	108.7	-0.8	889,257
2020-07-27 21:00:00	2020-07-27 23:00:00	35,000,000	0.008	0.003	4.38	1.10	0.13	11,120.3	-77.4	89.8	-1.2	1,104,155
2020-07-28 15:00:00	2020-07-28 19:00:00	27,000,000	0.006	0.004	3.48	1.11	0.07	11,172.8	-160.8	62.8	-0.4	480,057
2020-07-31 06:00:00	2020-07-31 08:00:00	18,000,000	0.005	0.004	2.04	1.17	0.08	11,152.5	-14.9	50.5	-3.4	313,457
2020-08-01 06:00:00	2020-08-01 08:00:00	15,000,000	0.008	0.005	3.68	1.34	0.09	11,648.1	34.3	97.7	2.9	1,329,122
2020-09-15 12:00:00	2020-09-15 14:00:00	19,000,000	0.005	0.004	2.96	1.43	0.67	10,914.9	-141.4	56.4	-0.4	458,476
2020-10-19 15:00:00	2020-10-19 17:00:00	10,000,000	0.006	0.003	5.38	1.37	0.67	11,804.9	-75.4	75.3	-1.0	817,733
2020-10-21 12:00:00	2020-10-21 16:00:00	29,000,000	0.003	0.003	3.53	1.43	0.63	12,423.1	417.5	43.0	0.1	583,482
2020-10-29 15:00:00	2020-10-29 17:00:00	15,000,000	0.005	0.004	2.93	1.50	0.56	13,440.6	148.8	65.1	0.4	444,188
2020-11-17 17:00:00	2020-11-17 23:00:00	60,200,000	0.008	0.008	4.97	1.79	0.47	17,795.5	-113.9	141.0	-1.2	1,857,423
2020-11-30 14:00:00	2020-11-30 16:00:00	26,000,000	0.014	0.008	4.82	1.55	0.52	19,494.6	-304.6	268.1	-0.9	6,741,367
2020-12-19 14:00:00	2020-12-19 16:00:00	27,000,000	0.009	0.007	3.17	1.49	0.58	23,304.2	679.2	202.2	0.3	2,945,704

Material USDT outflow events from the tether treasury towards exchanges and corresponding estimates of the equally-weighted market portfolio index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value based on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_{γ} , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.

Similarly, 11 USDT outflow events towards other unknown entities are considered material and included in Table 3.6. Again, the significant events are located primarily in Q2-Q4 of 2020; the average event duration is 2.7 hours, with most events lasting 2 hours and some cases of 3, 4 and 5-hour events. The total USDT sent to these unknown entities in the above events amounts to 350 million USDT, ranging between 5 million - 100 million. The estimated average abnormal return for the above events ranges between 0.52% - 1.61% with an average value of 0.97%. The regression adjusted R^2 ranges between 0.47 - 0.77 with an average value of 0.58. The estimated price effect ranges between \$49 - \$264, with an average of \$125, and the $P^{Effect}/\Delta P$

ratio again often exceeds 1 in magnitude and is often negative. Finally, the total estimated damage amounts to \$30 million, ranging between \$580k and \$6.6 million. Similar results are provided in Appendix B, in Tables B1 and B3 for the cap-weighted MVDA index model and in Tables B2 and B4 for the PCA-based model.

Table 3.6: Material USDT outflows from the tether treasury to unknown entities

Start	End	USDT	$\hat{\gamma}$	SQ	t_γ	SQ_{std}	Adj. R^2	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-05-13 17:00:00	2019-05-13 19:00:00	4,995,000	0.016	0.006	6.62	1.53	0.50	7,800.1	82.7	125.5	1.5	2,044,148
2020-04-23 14:00:00	2020-04-23 17:00:00	25,584,859	0.008	0.005	4.37	1.44	0.77	7,548.3	11.7	59.1	5.0	3,470,125
2020-04-29 20:00:00	2020-04-30 00:00:00	22,001,000	0.008	0.005	4.89	1.40	0.71	8,931.3	-115.1	70.7	-0.6	1,209,568
2020-05-07 17:00:00	2020-05-07 22:00:00	36,000,000	0.009	0.006	4.59	1.27	0.57	9,844.1	55.5	84.1	1.5	4,128,227
2020-06-01 23:00:00	2020-06-02 03:00:00	50,000,000	0.012	0.008	4.73	1.53	0.52	10,211.2	-116.3	119.1	-1.0	5,399,317
2020-07-22 22:00:00	2020-07-23 00:00:00	15,000,000	0.005	0.003	4.44	1.33	0.52	9,507.4	17.0	49.7	2.9	582,353
2020-10-21 14:00:00	2020-10-21 16:00:00	10,000,000	0.008	0.003	6.20	1.40	0.64	12,741.4	99.1	101.2	1.0	1,765,803
2020-10-21 22:00:00	2020-10-22 00:00:00	25,000,000	0.012	0.003	8.85	1.41	0.62	13,219.7	-298.0	156.7	-0.5	2,337,395
2020-11-20 16:00:00	2020-11-20 18:00:00	10,000,000	0.008	0.008	3.12	1.74	0.47	18,694.1	-110.7	152.1	-1.4	1,246,770
2020-11-30 14:00:00	2020-11-30 16:00:00	49,999,999	0.014	0.008	4.74	1.53	0.52	19,494.6	-304.6	264.3	-0.9	6,647,181
2020-12-19 17:00:00	2020-12-19 19:00:00	100,000,000	0.008	0.007	2.86	1.55	0.57	23,908.0	-96.4	190.4	-2.0	1,601,355

Material USDT outflow events from the tether treasury towards other unknown entities and corresponding estimates of the equally-weighted market portfolio index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_γ , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.

The results described above are conceptually consistent with the findings of Ante et al. (2021a), Ante et al. (2021b) and Griffin and Shams (2020), in that certain stablecoin blockchain transactions are indeed identified as having a significant effect on bitcoin’s price, often contributing towards a negative price trend reversal. Instead of attempting to determine if tether transactions on the whole have a significant effect on crypto asset prices, as attempted e.g. by Lyons and Viswanath-Natraj (2020) and Wei (2018) who find no significant aggregate effect, the methodology presented in this Chapter accounts for the possibility that only some of the USDT outflows from the tether treasury may cause significant positive abnormal bitcoin returns and therefore have a significant and potentially manipulative effect on bitcoin’s price. The results presented in this section focus on such significant USDT outflows as

indicated by the strictest materiality criteria, and provide actual estimates of the corresponding price effect and damages caused.

3.6 Summary and Conclusions

This Chapter has examined the topic of potential bitcoin market manipulation through large tether blockchain transactions originating from the centralized tether treasury. The methodology draws from securities fraud litigation studies to determine the materiality of price effects based on regression factor models and estimation of abnormal bitcoin returns during each USDT outflow event. The key contribution of this Chapter stems from the fact that damage estimates in securities fraud cases are not accepted in court unless an event study-based approach is employed, and the methodology developed is considered an improvement compared with similar research examined in the literature review. Additionally, significant detail is provided on the data retrieval process of blockchain transaction data, which are an uncommon data type and their data retrieval techniques are sparsely covered in the relevant literature.

Overall, the findings indicate that 30 out of a total 671 USDT outflow events originating from the tether treasury produce a significant positive abnormal return and therefore a material positive price impact indicative of price manipulation, based on the equally-weighted market portfolio index model specification, while results from alternate model specifications are also provided. The total damages from the above material events are estimated at \$54 million; by comparison, in the somewhat relevant [2021 legal case](#) where Bitfinex and Tether ‘deceived clients and market by overstating reserves and hiding approximately \$850 million in losses’, Tether agreed in an out-of-court settlement with the New York Attorney General’s office to pay a fine of \$18.5 million in penalties.

Regarding the limitations and potential extensions of the methodology presented here, it is important to note that intention to manipulate crypto asset prices cannot

be demonstrated with the methodology presented, and would instead be based on evidence extracted from investigative research similar to the article of Protos (2021*a*). Consequently, the estimated damages may be attributed to multiple offenders depending on the specifics of the legal case; similarly, in an actual legal case, damages would be estimated based on the actual traded volume transacted by the plaintiffs, instead of the approximation used in the methodology developed in this Chapter.

THE TOKENOMICS OF CROWDFUNDING

4.1	Introduction	133
4.2	Literature Review & Hypotheses Development	136
4.2.1	Literature Review	136
4.2.2	Hypotheses Development	142
4.3	Data	150
4.3.1	Variables	152
4.3.2	Sample Statistics	158
4.4	Models	161
4.5	Empirical Results	165
4.5.1	Sample Period 2017 – 2020	165
4.5.2	Sample Period 2021 – 2022	177
4.6	Summary and Conclusions	184

4.1 Introduction

This Chapter examines the fundraising success determinants of token offerings and their evolution between 2017 – early 2022, using linear and probit regression models. In this Section we first provide a brief introduction to token offerings and then a short discussion on our research hypotheses.

Token offerings are decentralized crowdfunding campaigns. A token offering is conducted by a start-up venture in order to raise funding for a blockchain- or crypto-asset-related project at an early stage of development. A digital token is issued using a smart contract blockchain platform such as Ethereum, or more recently a crypto asset exchange token launch platform such as the Binance Launchpad or a decentralized launchpad such as PancakeSwap.¹

All aspects of the token itself and the offering’s design are defined in smart contract code often written following commonly-used coding standards such as the ERC-20 token standard on the Ethereum platform. Ideally, the token serves an integral purpose in the venture’s ecosystem, which is very often to grant access to an offered product or service; if this is the case, the token is classified as a utility. More rarely, it may be that the token grants the holder cash-flows similar to stock dividends or bond coupon payments; such crypto assets are classified as security tokens and security token offerings (STO) are regulated more strictly in certain jurisdictions. In fact, the increase in regulatory oversight for security tokens has driven the majority of ventures to classify their token as a utility crypto asset. The primary difference of a token offering compared with an IPO is that usually a token does not afford its holders any ownership rights over a company (OECD, 2019).

For the duration of the token offering, a certain amount of token units (henceforth tokens) are offered for sale to potential investors. Investors can deposit crypto assets such as bitcoin (BTC) and ether (ETH) or on certain occasions fiat currencies to the venture’s digital wallets and accounts, to receive the venture’s tokens at a predefined

¹A digital token is a crypto asset that exists on a non-native blockchain. Note that in some cases the crypto asset offered to prospective investors is deployed on a native blockchain created by the venture.

rate of exchange. This rate of exchange can include various rebates for early investors (bonus scheme) or for advertising the offering on social media (bounty scheme). The venture often sets a minimum and maximum fundraising target, but does not have to do so. The minimum fundraising target, called a soft cap, defines a lower funding boundary for the offering: if the token offering does not raise more funding than the soft cap, it is usually cancelled and all funds are returned to investors. The maximum fundraising target, called a hard cap, defines the upper funding boundary: if the funds raised exceed the hard cap, the excess is often – but not necessarily, returned to investors.

For instance, [CoinPoker](#) is an Ethereum-based poker platform. Its token is called CHP (Chips) and is an ERC-20 token used on CoinPoker for all bets and payments made to and from a user’s account, eliminating the lack of transparency prevalent in traditional online poker platforms and simplifying deposits and withdrawals. A total of 275 million CHP tokens were created, 75% of which were offered for sale in the token offering. The CHP token sale took place between 19 – 29 January 2018 and investors could purchase CHP tokens at the rate of 4,161 CHP per ETH (with rebates offered to early investors), approximately equivalent to a dollar price of \$0.2497 per CHP token. CoinPoker set a low soft cap of 15 ETH – equivalent to \$15,500 at the launch of the offering, and a hard cap of 93,839 ETH, equivalent to \$100 million; the token sale raised a total amount of 52,600 ETH – equivalent at the time to \$55.5 million.

A report by OECD (2019) summarizes the differences and similarities between token offerings and more traditional funding channels such as IPOs, reward- and equity-based crowdfunding and venture capital. While IPOs are used by established companies with a revenue track record, token offerings are very often used by start-up ventures that are not even incorporated. Token sales and equity- and reward-based crowdfunding are quite similar funding mechanisms but there are some important differences, such as the existence of a central counterparty in crowdfunding versus the decentralized structure of token sales. However, this boundary is becoming less

clear in recent years with the emergence of token offering launchpad platforms on centralised and decentralized crypto asset exchanges. Finally, venture capital (VC) funding is considered complementary to token offerings, with VC funds participating in several private token sales preceding the public token offering and often providing ventures with expertise, networking and strategic advice.

Our research hypotheses focus on the specific characteristics of token offering fundraising success determinants between 2017 – early 2022. We proxy fundraising success with the amount of funding raised and also by whether the minimum funding required, the soft cap, is exceeded. We examine factors derived from the venture, token and offering characteristics, as well as additional common factors such as the average price level and momentum of ether during each offering and the launchpad platform, if used.

The sample used consists of 2,926 token offerings and covers the entire 2017 – early 2022 period, which – at the time of writing, constitutes one of the most comprehensive studies on the subject. We introduce explanatory variables that are previously omitted in the relevant literature, we discuss alternate perspectives on variables already used, and we revisit variables for which there are conflicting findings. Importantly, this study is the first relevant research to identify a revival in the market of token offerings in the 2021 – early 2022 period, with a significant shift in the fundraising success determinants.

The remainder of this Chapter is structured as follows: in Section 4.2 we review the relevant papers in the token offering success literature and develop hypotheses on the determinants of fundraising success; in Section 4.3 we discuss the data and the construction of our variables and provide their sample statistics and a correlation analysis; in Section 4.4 we present our regression model specifications, formalise the interpretation of the results and discuss the process of standardising some of the variables; Section 4.5 presents and discusses our empirical results based on regression models examined separately for the 2017 – 2020 and 2021 – early 2022 periods.

Finally, Section 4.6 summarizes our findings and provides suggestions for further research.

4.2 Literature Review & Hypotheses Development

The majority of the business and finance literature that examines token offerings focuses on the determinants of fundraising success and also post-offering token performance. In this chapter, we focus on the fundraising success strand of the literature. Note that the terminology used to describe token offerings in the relevant literature is varied, due to the lack of standardization and the ever-changing nature of the crypto asset space; authors very often refer to token offerings as ICOs (initial coin offerings), but given the emergence of several variants such as STOs (security token offerings), IEOs (initial exchange offerings) and more recently IDOs (initial decentralized exchange offerings), we choose the more generic term ‘token offering’.

4.2.1 Literature Review

A relevant search for published and conference papers on the Scopus database yields a total of 136 papers, 60% of which are published between 2020 – 2022.² In the following, we first discuss the theoretical framework of token offering fundraising success as well the findings in the relevant literature, followed by the development of hypotheses based on our primary research questions: ‘what are the determinants of token offering fundraising success’ and ‘how do these determinants change over time as the token offering space evolves’.

Table 4.1 provides a summary of the results for the relevant papers using the most common model setup in the token offering fundraising literature: a linear regression model using (log) amount raised as the dependent variable. For the commonly

²The full Scopus query used is: TITLE-ABS-KEY(("initial coin offering" OR "ICO" OR "token offering" OR "initial exchange offering" OR "IEO" OR "security token offering" OR "STO") AND ("success" or "fundraising" or "funding" or "crowdfunding" OR "performance")) AND (LIMIT-TO(SUBJAREA,"BUSI") OR LIMIT-TO (SUBJAREA,"ECON")) AND (LIMIT-TO(DOCTYPE,"ar") OR LIMIT-TO(DOCTYPE,"cp"))

used explanatory variable, we show whether a paper reports a positive, negative or insignificant effect, and also report each paper’s sample period, sample size and linear regression R^2 , using the adjusted value where available.

Table 4.1: Empirical results in the token offerings literature

Authors	Bourveau et al. (2022)	Campino et al. (2022)	Lyandres et al. (2022)	Belitski and Boreiko (2021)	Boreiko and Risteski (2021)	Czajka and Röder (2021)	Philippi et al. (2021)	Thies et al. (2021)	Yen et al. (2021)	Ante and Fiedler (2020)	Montaz (2020a)	Montaz (2020b)*	Rosenboom et al. (2020)**	Aggarwal et al. (2019)	Albrecht et al. (2020)	Boreiko and Vidusso (2019)	Chen (2019)	Fisch (2019)	Amsden and Schweizer (2018)	Blaseg (2018)*	Lee et al. (2018)*
Data start	4/14	5/16	1/13	1/13	1/13	7/14	1/14	11/16	1/17	4/17	8/15	8/15	8/15	1/15	1/17	1/13	1/15	3/16	1/15	1/14	1/16
Data end	2/18	5/20	11/18	12/17	10/17	1/18	10/20	3/19	1/18	10/18	7/18	4/18	12/17	9/18	3/18	9/17	3/18	3/18	3/18	12/17	5/18
N	200	428	980	166	272	664	357	1597	841	151	495	132	630	853	522	178	479	423	214	670	727
Price		+				o	o		o				o		o					o	
Supply											+		+					+	+		
Github	o	o		+	+		+						+		+			o	o		
Whitepaper			o		+					o						o					
Presale		o	+	+	o	–	–					–	+					o	o		–
Rating	+							+	+			+		+		+					+
Team size	+	+	+			o	+			+			+	+			+		+	o	
# Advisors						o	+			o								o			
ETH		+			o								+			o			o	+	
Bonus		+		o	o	–	o						o						o	o	–
Distribution			–			–	o						o				o	–	o	–	
Duration					–		o		–		–	–	o		–		–	–		o	
R^2	0.28	0.29	0.39	0.37	0.36	0.31	0.38	0.11	0.14	0.38	0.23	0.18	0.23	0.18	0.37	0.30	0.22	0.42	0.28	0.14	0.03

Empirical results in the relevant literature. Only results on linear regression models are included, with log amount raised as the dependent variable. The first row lists the relevant papers; papers denoted with * use the actual amount raised instead; ** denotes the use of $\log(1 + \text{Raised})$. Note that several papers present results on multiple model specifications; whenever this is the case, we only include the most informative one. The data sample period and size (N) used in each paper are shown in rows 1-3; the final row shows the regression R^2 , using the adjusted value when available. The remaining rows exhibit the results on relevant explanatory variables; a + (–) sign denotes a positively (negatively) significant coefficient and an ‘o’ denotes insignificant coefficients. Note that we combine the results on all variants of the token offering rating variable (Experts Rating, Team Rating, Vision Rating, Project Rating) into one representative Rating variable, similar to Bourveau et al. (2022).

The sample periods examined in the literature commonly span 2014 – 2018, with only Campino et al. (2022) and Philippi et al. (2021) extending their sample into 2020, and none of the papers examine token offerings in 2021. While a small number of offerings takes place prior to 2017, e.g. Philippi et al. (2021) report approximately 70

such cases, the number of token offerings rises significantly in mid-2017 and onwards; despite this observation, several recently published papers shown in Table 4.1 do not extend their sample period beyond 2017. The sample sizes examined vary significantly between approximately 200 – 1,600 token offerings due to the use of alternate and complementary data sources and the limited availability of data depending on the variables examined; for instance, Lyandres et al. (2022) study the fundraising success of 980 token offerings between 2013 – 2018 and also perform an extensive review of data quality and availability on multiple online sources. We further note the existence of the ‘Token Offerings Research Database’ (TORD) by Momtaz (2021), a very promising attempt at creating a ‘one-stop’ data source for token offerings, covering a variety of venture and offering characteristics on approximately 6,000 offerings with final sample size depending on data availability.

All papers included in Table 4.1 proxy token offering fundraising success by the amount of funding raised in the offering, and almost all use a log transformation. Other fundraising success proxies are also used in the literature such as: the amount funding raised relative to the hard cap, as done e.g. by Roosenboom et al. (2020) in an alternate specification and also by Lyandres et al. (2022); whether the soft or hard cap are exceeded, used for instance by Sharma and Zhu (2020), Giudici and Adhami (2019) and Roosenboom et al. (2020); the square root of the amount raised used by Burns and Moro (2018); whether the amount raised exceeds zero (Blaseg, 2018) or a specific amount such as \$10,000 (Lyandres et al., 2022) or \$0.5 million (Lee et al., 2018).³ Linear or non-linear (logit and probit) regression models are used to explain the above success proxies, depending on whether the dependent variable is continuous or binary. The commonly-used explanatory variables are derived using the characteristics of the venture, the token offering and the token itself, the venture’s social media activity, venture ratings from dedicated rating and ranking websites and also market and sentiment factors.

³Post-offering performance proxies include: whether the token is trading on crypto asset exchanges; the token’s return on the first day of trading and others.

The theoretical framework of token offering fundraising success is based on similar frameworks developed in the entrepreneurial finance literature. For instance, Thies et al. (2021) suggest that the components of a token offering exhibit the typical characteristics of crowdfunding and IPOs. Ahlers et al. (2015) examine funding success on equity crowdfunding platforms and develop a framework in which the key factors of a venture's funding success are: (i) observable characteristics of venture quality and (ii) reduced investor uncertainty. As investors are looking to maximize their future returns, they will prefer to invest in high-quality ventures which are more likely to provide higher returns. Ahlers et al. (2015) suggest that venture quality is demonstrated to prospective investors via signals of human, social and intellectual capital, and Baum and Silverman (2004) also present a similar argument for startups that receive venture capital funding. Connelly et al. (2011) further argue that an observable signal of venture quality also needs to be costly to produce, so prospective investors are likely to disregard cheap signals as 'cheap talk'. Parallel to the concept of venture quality signals, reduction of investor uncertainty relating to the proposed venture is also key: if prospective investors are not in possession of sufficient and accurate information regarding a proposed venture, their assessment of it is restricted and therefore they are less likely to invest. For instance, Ahlers et al. (2015) suggest that a higher equity share retained by the venture launchers and the availability of financial projections can reduce uncertainty and mitigate the risk of asymmetric information.

In the following, we examine adaptations and extensions of the above theoretical framework, for the case of token offerings. For instance, Bourveau et al. (2022) and Blaseg (2018) focus on the role of information disclosure in reducing investor uncertainty and driving token offering fundraising success. Czaja and Röder (2021) extend the framework of Ahlers et al. (2015), viewing the level of investor familiarity and attention relating to a proposed venture as a separate factor contributing to higher funding amounts in token offerings. Venture quality signals are also examined extensively in the context of token offerings. Philippi et al. (2021) suggest that

technological, venture, and campaign characteristics act as signals of the venture’s technological capabilities and Fisch (2019) proxies technological capability using the existence of a technical whitepaper mentioning a patent and high Github activity, suggesting that such indicators can serve as signals that reduce information asymmetry.

Campino et al. (2022) propose that token offering projects can be viewed as ventures operating in an ‘open systems model’, similar to crowdfunding ventures. In this context, Campino et al. (2022) argue that providing prospective investors with higher levels of quality information in the open systems relationships contribute positively to the funding outcome of token offerings. Chen (2019) proposes that highly credible and easily-interpretable signals have a positive effect on fundraising success, similar to the proposal of Connelly et al. (2011). Thies et al. (2021) also focus on signalling theory and distinguish between endogenous and exogenous signals, i.e. signals that can or cannot be influenced directly by the emitter; they suggest that both endogenous signals such as the choice of social media channels and exogenous signals such as independent expert ratings can influence token offering funding success and also moderate the effect of hype as proxied by the price of Bitcoin. Yen et al. (2021) and Aggarwal et al. (2019) also suggest that external ratings strengthen the credibility of information provided to prospective investors and Belitski and Boreiko (2021) further interpret venture capital support, whitepaper publication and presence on Github as venture quality signals that have a positive effect on fundraising success.

However, several authors caution of a potential negative aspect in the extended use of venture quality signals to attract prospective investors. Ante and Fiedler (2020) follow the concept of ‘cheap talk’ signals discussed by Ahlers et al. (2015) and Connelly et al. (2011), somewhat similar to the endogenous-exogenous signal distinction made by Thies et al. (2021); Ante and Fiedler (2020) suggest that signals under the venture’s direct influence such as its website or social media channels may indicate biased or potentially faked signals, while they argue that the use of false signals by independent external parties seems less probable. Momtaz (2020a) focuses

on moral hazard in signalling and suggests that token issuers could be incentivised to bias signals of venture quality to their advantage, given the attraction of investors to positive signals and the lack of institutional agents that detect and punish biased signals; as a result, the token offerings of ventures that do not send biased signals could in fact be less successful. Boreiko and Vidusso (2019) examine the token offering ratings assigned by independent websites, suggesting that they could serve as valid venture quality signals but also warn of moral hazard in the issuer-pays model used by aggregator websites for token offering information.

Additionally, several authors examine venture and token offering characteristics as fundraising success determinants that are not assigned to or only loosely related to the previously discussed factors of venture quality characteristics, investor uncertainty reduction and investor familiarity and attention. For instance, Roosenboom et al. (2020) suggest that the use of bonus schemes and a longer planned offering duration have a negative influence on fundraising success, in the sense that the venture issuers display a lack of self-efficacy to potential investors. Amsden and Schweizer (2018) consider the ETH/USD traded price and propose that when it is high, investors may not wish to part with their ether holdings and be less incentivised to participate in token offerings, while the volatility of the return on ether's dollar price can be an indicator of increased regulatory uncertainty about crypto asset investments in general and in such cases, investors hasten to participate in an offering for fear of missing out on a limited-time opportunity. Albrecht et al. (2020) propose that the initial price setting of tokens has no major impact on the financial success of an offering, as promising tokens could sell at the same price as underwhelming tokens, and also suggest that online search trends regarding blockchain technology and crypto assets are positively linked to fundraising success, analogous to the concept of bandwagon effects, in which trends spread on a large scale, build positive feedback loops, and pressure individuals to conformity.

Momtaz (2020b) argues that management team quality and the quality of the venture's profile have a positive effect on fundraising success, being at the core of

principal-agent models; in the absence of effective corporate governance mechanisms, poor managerial quality translates directly into agency costs and deters prospective investors. Giudici and Adhami (2019) propose that attributing a legal jurisdiction to the venture’s project improves fundraising success, serving as a lever for the enforcement of investor rights. Giudici and Adhami (2019) further suggest that an increased venture team size can attract prospective investors and increase funding success, as more individuals share in the reputational risk of joining the venture. Similarly, Blaseg (2018) also argues that increased disclosure about the commitment of the venture team, as proxied by its size, positively affects funding success.

As described above, token offering fundraising success factors can be categorized into various theoretical frameworks, and multiple and sometimes conflicting arguments are presented as to the potential effect of these factors. This is also reflected in the empirical results of the relevant literature, shown in Table 4.1. For example, the Distribution variable reflects the share of total token supply offered to investors, so a lower share implies higher equity retention by the venture, which in turn should reduce investor uncertainty as suggested by Ahlers et al. (2015). However, only half of the papers that examine this variable find a significant negative effect that would be consistent with the above argument. Similarly, relative consensus is reached for several variables, such as the negative effect of longer offering duration on the amount raised, but even then there is often at least one finding of an insignificant or opposite effect; for instance, Philippi et al. (2021), Roosenboom et al. (2020) and Blaseg (2018) find that duration has no significant effect on the amount raised. Notable exceptions to this are the rating variable, for which all papers find a significant positive effect, and also the Supply variable which is used less extensively, possibly due to its high negative correlation with the token’s initial price.

4.2.2 Hypotheses Development

Having examined the theoretical frameworks, hypotheses and findings in the relevant literature, we now discuss the incremental contributions introduced in this chapter,

followed by a presentation of our hypotheses on the determinants of token offering fundraising success.

We employ the amount of funding raised and also – where possible, the soft cap exceedance as success proxies. We select possible success determinants from the characteristics of each token offering, the token’s technical structure, the venture’s online and social media presence and also indicators relating to the overall token offerings ecosystem and crypto asset market; we introduce explanatory variables that are not used in the relevant literature, such as the combination of token offering supply and offered price to form the token’s target market capitalization, as well as the launchpad platform used by token offerings in 2021 – early 2022; we examine explanatory variables in which there is lack of consensus amongst the relevant literature – such as the token presale, and also offer new perspectives on existing variables such as the token offering rating and tax-haven domicile.

We examine the entire 2017 – early 2022 period, which – at the time of writing, constitutes the most comprehensive study on the subject in terms of sample period coverage. We confirm and extend the findings in the relevant literature for the 2017 – 2019 period; we examine separately the evolution of success determinants during the earlier 2017 – mid-2018 boom period, in the subsequent period of reduced investor interest in token offerings, and we also identify the near total stand-still of the token offering space during 2020 due to the Covid-19 pandemic. Importantly, this study is the first to identify 2021 – early 2022 as a period of renewed investor interest in token offerings, with December 2021 and January 2022 exhibiting new all-time highs in the number of token offerings completed per month. We also identify a significant shift in the token offering market with the emergence and near-complete domination of initial decentralized exchange offerings (IDOs), as well as a shift in the fundraising success determinants, with the launchpad platform used apparently playing a vital role in explaining the fundraising success of token offerings in 2021 – early 2022.

In the following, we first present the statement of each hypothesis, followed by a brief discussion of the rationale and whether similar hypotheses are proposed in the relevant literature.

H1: The fundraising success of token offerings increases with the target market cap and decreases with the soft cap level.

We define the target market cap in an offering as the product of the token's price and its total supply, and suggest that it provides an equal basis of comparison between tokens for investors, leading to a reduction in investor uncertainty. Price and Supply are used independently in the literature, e.g. Roosenboom et al. (2020) use both the token's initial price and the token's total supply as separate determinants in their regression analysis, and similarly, Albrecht et al. (2020) suggest that the token's price in the offering has no significant impact on fundraising success; however, to the best of our knowledge, no other papers in the relevant literature combine Price and Supply in a single variable. A token's initial price and its total supply are inversely related: by setting the token's price and its total supply the venture is actually setting the target for the initial market capitalization of its token, calculated as the product of total token supply and price; if a venture sells its tokens at a very low price, it needs to sell a large number of them to raise sufficient funds, and vice versa. We therefore propose that it is the combination of price and supply information that allows comparisons between token offerings, reducing uncertainty and ultimately driving the decision to invest in a token.

We now discuss the role of the soft cap level. The soft cap in an offering represents the minimum funding required, ideally estimated by the venture based on the expenses of creating a minimum viable product, including development and operational expenses, wages, marketing costs and perhaps a provision for other unexpected costs. Setting the soft cap too high can instil doubt to prospective investors who may wait until the offering attracts enough funding before investing

themselves. Moreover, ventures sometimes decide to lower their soft cap and instead use funds of their own or raised via a token presale to balance the funding gap; this may attract more investors as the venture is seen to have more ‘skin-in-the-game’, again reducing investor uncertainty. For instance, Roosenboom et al. (2020) use the soft cap in a logit regression model and find that it has a negative effect on the probability of soft cap exceedance. Note that a lower soft cap may also have a positive effect on the amount raised but we only use this variable to explain the probability of soft cap exceedance, because several offerings in our sample are uncapped.

H2: The fundraising success of token offerings increases with signals of public credibility, but only if they are genuine.

Several token offering and venture characteristics can be viewed as signals that are used to reduce the asymmetry of information between a venture and prospective investors. As mentioned previously, the use of signalling and asymmetric information theory in explaining token offering fundraising success is extended from the crowdfunding literature, based on much earlier studies by Akerlof (1970), Spence (1973), Leland and Pyle (1977) and Spence (2002). The more recent studies of Baum and Silverman (2004), Hsu and Ziedonis (2013) and Ahlers et al. (2015) emphasize that crowdfunding ventures need to signal their technological capabilities and intellectual capital effectively in order to reduce information asymmetry between investors and the venture. In order to do so, ventures need to use signals that are observable by investors and are costly to realize and imitate. High-quality ventures therefore attract more funding if potential investors are able to distinguish them from ventures of lower quality (Connelly et al., 2011). Extending from this literature, we examine token offerings and distinguish between information availability and informational content to characterize available signals as genuine or ‘cheap’ and examine whether investors are sensitive to this distinction. For instance, the online availability of a

venture's source code or whitepaper are 'cheap' signals while an assessment of the source code or whitepaper's quality is a genuine one.

A venture launching a token offering can signal its quality to prospective investors in several ways: a positive rating on independent rating and ranking websites enhances the venture's credibility, as such websites are often visited by prospective investors. A successful token presale can serve as a proof-of-concept and provide evidence of investor interest and belief in the venture's success. A whitepaper provides detailed general and technical information about the proposed project, while an active presence on Github can help disseminate similar information on an on-going basis. The public credibility and competence of the venture's team members and advisors, proxied by the size of these teams, can also have a significant effect in the success of a venture's token offering.

As discussed previously, Boreiko and Vidusso (2019) warn about certain caveats in token offering ratings, such as the issuer-pays model used by rating and ranking aggregator websites such as ICObench. Interestingly, Hartmann (2018) in a Medium article actually provides evidence that token issuers can purchase favourable ratings for their offering! These fake ratings are purchased from unknown entities who then post the favourable reviews on ICObench; it is even possible to dictate the text content of such 'expert reviews' on ICObench. While ICObench claims to apply stringent vetting of third-party experts registered on the website, Hartmann (2018) suggests that 'ICObench and other rating platforms at least approve of or tolerate these dubious practices on their platforms'. Therefore, ratings may not have a significant effect on fundraising success, as such revelations perhaps discourage investors from trusting them.

Furthermore, hypotheses on signals of public credibility are quite common in the relevant literature but the findings often show that such variables are insignificant in explaining fundraising success, as shown in Table 4.1. Moreover, Momtaz (2020a) warns of a bias caused by exaggerated or false signals and Ante and Fiedler (2020) suggest that cheap signals should not be trusted by investors, although this is not

always found to be the case. We believe that these warnings have merit, and wish to test them on a more recent sample of token offerings. For instance, the availability of a Github code repository and a whitepaper can be argued to be cheap signals, as they only represent the availability of information and not its content.

Additionally, the overwhelming majority of token offerings in 2021 – early 2022 are conducted on centralized (IEOs) and decentralized (IDOs) exchange launchpads such as Binance, Pancakeswap and Polkastarter, and these launchpad platforms are quickly becoming key intermediaries in the token offerings ecosystem. Token offering launchpad platforms are therefore becoming an important factor in venture fundraising success, to the point that prospective investors perhaps consider offerings conducted on well-known launchpad platforms as more credible.

H3: The fundraising success of token offerings increases with investor belief in the crypto asset class, as well as investor attention and hype.

It is an established observation that large cap crypto assets, particularly bitcoin and ether often drive the entire asset class. The dollar price of bitcoin and ether and their return and/or momentum are often used in the relevant literature as control variables; for instance, Amsden and Schweizer (2018) suggest that when the dollar price of bitcoin and ether are higher, investors prefer to maintain their bitcoin and ether holdings and are therefore less willing to invest them in token offerings. Thies et al. (2021) also consider the hype effect of bitcoin’s price and propose that higher fundraising success is caused by investor herding due to hype. We propose additionally that the dollar price of ether acts as proxy for investor confidence in the entire crypto asset class. Investors are therefore more likely to participate in token offerings when the prices of bitcoin and ether are higher.

Increased hype may also attract prospective investors, and ventures that are active on Twitter have the ability to engage with the investor community and attract larger amounts of funding; alternatively, if investors are unaffected by Twitter-related hype,

then the availability of a venture’s Twitter profile should have no effect on funding success. Finally, in order to examine the effect of investor attention around crypto assets on offering fundraising success, we further introduce the Google search index based on the number of Google web searches, Google news searches and YouTube searches on the terms ‘Blockchain’, ‘Bitcoin’ and ‘Ethereum’.

H4: The fundraising success of token offerings is higher for ventures domiciled in a tax-haven or with an undisclosed domicile.

Given the immensely speculative nature of the crypto token market, investors are likely to worry less about token holder rights, legal recourse and dispute resolution and instead prefer increased versatility to invest in and trade crypto tokens and more favourable taxation schemes. As a result, they may prefer to invest in ventures registered in offshore financial centres (OFC) i.e. tax-havens or with an unspecified domicile, in order to avoid strict regulations on crypto tokens’ security classification, taxation and trading. Despite this argument, the findings in the relevant literature, e.g. by Amsden and Schweizer (2018), show that a tax-haven domicile has no significant effect on funding success and Huang et al. (2020) further find that tax-haven jurisdictions do not attract a larger number of token offerings.

Additionally, as more token offerings are conducted in launchpad platforms of centralized and decentralized exchanges, offering investors the ability to trade their tokens immediately on the exchange platform, it is possible that investors are more interested on the regulatory status and domicile of such integrated token launchpad/exchange platforms, in which case the domicile of individual ventures is less important.

H5: The fundraising success of token offerings are unaffected by the existence of a bonus scheme, the token's distribution and the duration of the offering.

A bonus scheme can give the impression of a venture desperate to attract funding and may deter investors; however, if no other deterring factors are present, the price rebates implied in the bonus scheme may in fact attract more investors. For instance, the Ethereum offering included an implicit bonus scheme in 2014; for the first 14 days of the token sale, 1 BTC could buy 2,000 ETH and subsequently the price increased linearly so that on the final day of the sale 1 BTC could buy 1,337 ETH. There is no consensus about the effect of a bonus scheme in the relevant literature; for instance Roosenboom et al. (2020) propose that bonus schemes negatively impact fundraising success while Adhami et al. (2018) suggest the exact opposite. Token distribution refers to the share of the token's total supply that is offered to investors in the public offering.

A lower token distribution means that the venture retains a significant share of the tokens and therefore has more 'skin-in-the-game', consistent with the suggestion of Ahlers et al. (2015) that a higher equity share retention by the venture reduces investor uncertainty. However, if the share of tokens retained by the team is very high, the project may be tempted to make a quick profit by 'dumping' its tokens on the market once the token begins to trade on exchanges.

A very long offering duration may indicate that the venture is not confident of raising the required amount within a short period, but on the other hand immensely successful offerings such as that of EOS have lasted more than one year. A consulting report by a Swiss law firm (Baker McKenzie, 2018) suggests that a realistic time horizon for a legally compliant token offering is between 9 and 12 months, and the token offerings in our sample have much lower average duration but with a lot of dispersion within the sample. Again, as discussed previously, there is partial

consensus in the relative literature that duration has a significant negative effect on fundraising success.

4.3 Data

In this Section we discuss the data sources used to obtain our sample of token offerings, together with an overview of data quality in the token offering space. We then define the dependent and independent variables used in the regression models and provide preliminary analysis in the form of sample statistics and a correlation analysis.

We treat [ICObench](#) and [Cryptorank](#) as our primary data sources and create Python-based web scrapers to collect a data sample of 6,613 offerings that take place between January 2015 – January 2022. There are however significant gaps in several variables in this dataset that limit our sample to 2,926 offerings: specifically, the sample consists of 1,607 token offerings for the more recent January 2021 – January 2022 period; similarly for the January 2017 – December 2020 period, the final sample consists of 1,319 offerings when examining the amount raised as a success proxy and to 902 when we proxy success with soft cap exceedance, despite considerable efforts made to fill in missing/unknown values from other ranking/rating websites such as [ICOdata](#), [TokenData](#), [ICOdrops](#), [ICOrating](#) and [Neironix](#). Obtaining reliable data on offerings ending in 2020 is especially difficult and we often resorted to manual data collection from the above sources and launchpad websites such as [Coinlist](#) and also from the following ranking and rating websites: [Bestcoinlist](#), [Cryptototem](#), [Coincurb](#), [Coinpaprka](#) and [ICOMarks](#).⁴

As mentioned previously, Lyandres et al. (2022) perform an extensive review of data availability and quality on token offerings, including several of the data sources mentioned above and Momtaz (2021) provides an alternative source in the token offerings research database (TORD). For instance, and similar to the

⁴We sporadically use the following sources as well: [ICOholder](#), [Cryptocompare](#), [Smith and Crown](#), [ICOMarketdata](#), [ICOstats](#), [Coincodex](#) and [Cryptodiffer](#).

review of Lyandres et al. (2022), the TORD database of Momtaz (2021) (v3 at the time of writing) provides available data on the amount raised for approximately 2,100 offerings and the overall supply for approximately 4,400 offerings. Note that depending on additional data requirements, these samples may be smaller. For example, compared with the Bancor token offering data examined by Lyandres et al. (2022), the TORD database is consistent with the consensus value on the amount raised (\$153 million) but does not include data on the hard cap, total token supply or number of tokens for sale, similar to many online sources such as ICObench.

Given that our analysis requires extensive coverage of the 2021 – early 2022 period, we choose not to use the TORD database and instead rely on our own data collection process described above. Even so, the difficulty in obtaining a complete dataset of token offerings persists and is also apparent when examining the availability of data on the 10 largest offerings by funding raised: EOS, LEO, Telegram Open Network, Dragon Coin, Huobi, Hdac, Filecoin, Tezos, Sirin Labs and Bancor. Obtaining all relevant data is only possible for Dragon Coin and Filecoin at the time of writing; for the others it is only possible to obtain partial information, even when examining alternate data sources. Moreover, we deliberately exclude Telegram Open Network and LEO from our sample. Telegram Open Network does not strictly qualify as a public token offering as it raised all of its funding during the presale and cancelled the public offering. We also consider LEO as an outlier given some strange circumstances surrounding its token offering such as the extremely hurried announcement of the offering, the significant lack of technical detail until after the token sale was concluded and the numerous legal issues and overall reputation of the Bitfinex exchange that launched the LEO token sale.

Note that we use a somewhat different set of variables depending on the sample period examined, 2017 – 2020 and 2021 – early 2022; this is because the characteristics of token offerings are quite different in 2021, requiring an alternate model configuration to explain funding success. The need for the above distinction is apparent when examining the phases in the token offerings market since 2017: an initial boom

period, a subsequent cool-off period, the pandemic period of 2020 and finally a period of renewed interest in token offerings. The boundary between the first two periods appears to be in Q2-2018 and this also means that the regression results for this part of our sample are more directly comparable with findings in the relevant literature, as most papers end their sample period in the second or third quarter of 2018. We choose to treat offerings conducted during 2020 separately, as we expect that the outbreak of the Covid-19 pandemic has a significant impact on crowdfunding via token offerings, and the 2021 – early 2022 period of renewed interest is also examined separately. The above distinctions are also visible in Figure 4.1.


In the upper panel of Figure 4.1, the number of token offerings that end in each month shows a decline in June 2018; starting in January 2020 there is another significant decline, concurrent with the worldwide decline in economic activity due to the pandemic; this decline shows some improvement towards the end of 2020 and the number of offerings peaks to new all-time highs within 2021. Similarly, as shown in the lower panel of Figure 4.1, the total amount raised by token offerings in each month confirms that the initial boom period seems to end in June 2018; the decline in interest towards the beginning of 2020 is even more apparent in the funding raised. Additionally, the funding raised within 2021 is lower compared with the 2018 boom period, but given the larger-than-ever volume of token offerings, this could indicate that token issuers are now driven less by greed and more by realistic budgeting.

4.3.1 Variables

Having provided an overview of available data sources on token offerings and our own data collection process, we now discuss the dependent and independent variables used in the regression models to test our hypotheses.

The dependent variables used in our regression models are the log of the amount raised measured in USD (**logRaised**) and a binary variable indicating whether the amount raised exceeds the offering's soft cap. These are the most commonly-used

Figure 4.1: Token offerings and funding raised between 2017 – 2022



Number of token offerings that end (red, upper panel) and combined total amount raised (green, lower panel) in each month, between 1 January 2017 - 31 January 2022.

proxies for token offering fundraising success in the relevant literature as described in Section 4.2.

These two funding success proxies are also complementary: using the amount raised as a dependent variable in a linear regression restricts the sample in that we need to exclude offerings with zero funding raised, to avoid biasing the results; conversely, we can include such cases in the sample when we restrict it to capped offerings and use the soft cap exceedance probability as the dependent variable in a probit model.

Using the amount raised as a dependent variable is reasonable because more funding can help a venture cover project costs which often increase due to delays or other unforeseen circumstances. On the other hand, soft cap exceedance provides a clear success-failure boundary for the offering; if an offering fails to reach its soft cap, very often the token sale is cancelled and all funds are returned to investors. Note that the soft cap exceedance proxy is used in the 2017 – 2020 dataset but not

in the 2021 – early 2022 period, as none of the token offerings in 2021 include a soft cap. We further choose not to use the hard cap in a similar variable; no adverse consequence is incurred if an offering does not achieve its hard cap.⁵

The explanatory variables used in our regression models are `logtCap`, `logSoftcap`, `Rating`, `Team`, `Advisors`, `Presale`, `Github`, `Twitter`, `Whitepaper`, `ETH price`, `ETH mom`, `OFC`, `Bonus`, `Distribution`, `Distribution-H`, `Google`, `logDuration`, together with additional control variables. We define the log of the target market cap **logtCap** as $\text{logtCap} = \log(\text{Price} \times \text{Supply})$, where `Price` is the token’s initial offered price and `Supply` is the total number of token units created. We define the variable in log form to reduce its kurtosis and produce a (closer to) linear relationship with the dependent variable `logRaised`. Similarly, the **Softcap** variable has a sample kurtosis well above 100, so we use **logSoftcap** instead.

A venture’s rating on ICObench incorporates a variety of components obtained directly from, or fact-checked by, independent third parties, rendering it a genuine credibility signal. **Rating** is the average rating assigned to each project in ICObench, which is a combination of two components: a rating assigned by ICObench based on project and token offering characteristics such as transparency and team quality;⁶ a rating based on the team, vision and product of each project, assigned by independent experts registered on the ICObench website.⁷ The rating ranges from 1 to 5 and allows for a single decimal digit. Note that several offerings ending in 2020 are not listed on ICObench so we obtain similar ratings from rating websites, as described above;⁸ note also that no ratings are available for token offerings in 2021 – 2022, again indicating reduced investor interest in such metrics. **Team** and **Advisors** are integer variables indicating the number of members in the venture’s team and its advisory board.

⁵Due to the lack of regulation and standardization in the token offerings space, there is no unified standard of conducting a token offering. There is no obligation to return all funds if the soft cap is not reached or to return extra funds if the hard cap is exceeded; it is indeed common practice, but projects can choose not to follow it.

⁶See the relevant [methodology](#).

⁷See the [information](#) on ICObench experts.

⁸Such ratings often range between 1 - 10, see e.g. the [ICOMarks rating methodology](#); where this is the case, we rescale so that all ratings included in our sample range from 1 - 5.

Presale is a binary variable indicating whether a venture conducted a preliminary token sale prior to the main offering. **Github** is a binary variable indicating whether the venture’s Github profile was active at the beginning of the offering, a proxy for the venture’s activity on Github. Each venture’s Github profile url is obtained from the relevant page on ICObench or similar websites; the date of the profile’s first Github commit is retrieved and compared with the offering’s start date to obtain the binary variable. Similarly, **Twitter** is a binary variable indicating whether the venture was active on Twitter at the beginning of the offering. It is constructed in a similar manner by retrieving the date of creation for each venture’s Twitter profile. **Whitepaper** is a binary variable indicating whether a venture’s whitepaper is available on ICObench.

ETH price and **ETH mom** refer to ether’s dollar price and 30-day momentum, i.e. the 30-day moving average of the returns multiplied by 100 to express in percentage points for scaling purposes, averaged within the duration of each offering. **Bonus** is a binary variable showing whether an offering includes a bonus scheme rewarding early investors with rebates in the price paid for purchasing the token. **Google** is based on the average number of combined searches averaged across Google, Google News and YouTube on any of the terms ‘Bitcoin’, ‘Ethereum’ and ‘Blockchain’. **Distribution** is the percentage of the token’s total supply offered for sale; while it is a continuous variable, there is significant clustering around numbers ending in 0 and 5 (e.g. 65%). Additionally, for the 2021 – 2022 sample there is more information available on all the shares of token supply for each venture, distributed e.g. to investors via a token sale, retained by the development team, retained for market making and/or liquidity provision and also airdropped for promotional or other purposes; the above information is combined in the **Distribution-H** variable based on the normalized version of Shannon’s H index as defined by Boydston et al. (2014). **Offerings** denotes the number of token offerings conducted by the same venture, and is only used in the 2021– 2022 sample as it is becoming a common characteristic

of token offerings. We also define **logDuration** as the log of the offering’s duration in days to reduce the variable’s kurtosis.

OFC (Offshore Financial Center) is a binary variable that indicates if a venture is domiciled in a tax-haven or has no registered domicile; Garcia-Bernardo et al. (2017) define ‘sink’ OFCs as countries with small domestic economies that attract and retain significant foreign capital and note that nearly all ‘sink’ OFCs are characterized as tax-havens. Our list of tax-havens includes the following countries: Anguilla, Bahamas, Belize, Bermuda, British Virgin Islands, Cayman Islands, Costa Rica, Curacao, Cyprus, Gibraltar, Guyana, Hong Kong, Isle of Man, Jersey, Liberia, Liechtenstein, Luxembourg, Malta, Marshall Islands, Mauritius, Monaco, Nauru, Panama, Saint Kitts and Nevis, Samoa, Seychelles, Singapore, St. Vincent and the Grenadines and Taiwan. We note that Singapore is not classified as a ‘sink’ OFC by Garcia-Bernardo et al. (2017) but we choose to include it in this list; Singapore is included by Zoromé (2007) in the list of countries that provide financial services primarily to non-residents, it ranks as the third-most dominant tax-haven globally according to Alstadsæter et al. (2018) and it also ranks very high in the [Financial Secrecy Index](#) which ranks jurisdictions according to their secrecy and scale of offshore financial activities; we pay particular attention to Singapore, as it is the most popular venture domicile in our sample.

The additional control variables included in our regression models are as follows: **Accepts BTC**, **Accepts ETH**, **Accepts other** and **Accepts fiat** indicate respectively the accepted payment methods in each offering: bitcoin, ether, other crypto assets and fiat money; quite reasonably, nearly all offerings accept ether since nearly all tokens are deployed as smart contracts on the Ethereum blockchain. **Capped** is a binary variable indicating whether the token offering includes a soft and/or a hard cap. **Ethereum-based** is a binary variable indicating that the token is launched on the Ethereum platform following an ERC token standard such as ERC 20. **KYC** and **Whitelist** are binary variables indicating whether any investor KYC (know-your-customer) or pre-registration need to be completed prior to participation

in the offering. **IEO** (initial exchange offering) is a binary variable used in the 2017 – 2020 sample period that indicates whether the token sale was launched on a centralised token launchpad, as shown e.g. on the ICObench [IEO list](#); for the 2021 – 2022 sample, we construct the **Type** control variable, indicating whether the offering is an ICO, IEO or IDO, with the majority of offerings categorized as IDOs, i.e. launched on decentralized exchange launchpad platforms.

We further construct industry-based control variables (**High**, **Medium** and **Low**) using ventures' categories as reported on ICObench for the 2017 – 2020 sample. Assuming the level of investment needed for the development of projects in each category, we form three groups representing ventures more likely to be invested in by high-, medium- and/or low-income investors:

- **High-income:** Artificial Intelligence, Banking, Big Data, Business services, DeFi, Electronics, Energy, Investment, Manufacturing;
- **Medium-income:** Art, Charity, Communication, Education, Fashion, Health, Infrastructure, Internet, Legal, Media, Real estate, Retail, Software, Tourism;
- **Low-income:** Casino & Gambling, Entertainment, Sports, Virtual Reality.

High-income investors are likely to be venture capital and angel investor funds, medium-income refers to retail investors and low-income refers to token purchases for casual uses such as betting, gaming etc. Note that the overwhelming majority of ventures list multiple categories as relevant, so there are overlaps between High, Medium and Low and the variables do not have a unit sum for each offering, which also avoids the ‘dummy variable trap’ of perfect collinearity between the constant term and a sum of binary variables. We also note that it is even possible and perhaps desirable for ventures to belong to all three categories as they are more likely to attract funding from a larger pool of investors; however, this only occurs for 3% of offerings in the 2017 – 2020 sample.

Finally, for the 2021 – 2022 sample, we control for each venture category separately using the **Category** variable and also examine the **Platform** control which denotes the launchpad platform used in each offering, if any.

4.3.2 Sample Statistics

We now discuss the sample statistics of key variables, separately for the 2017 – 2020 and 2021 – early 2022 periods.

Sample Period 2017 – 2020

We first examine the sample statistics and correlations for continuous and other non-binary variables: the dependent variable logRaised and the explanatory variables logtCap, Rating, Team, Advisors, logDuration, Distribution, ETH price and ETH mom.⁹ We also examine the correlations between these variables and again exclude binary variables since their distribution includes only two possible values, causing problems in interpreting the correlation coefficient.

In Table 4.2 we examine the sample statistics. The kurtosis of Raised and tCap, i.e. before applying a log, is well more than 100 indicating that the variables contain significant outliers; as shown in Table 4.2, the kurtosis of logRaised and logtCap is at more reasonable values. Similarly, we use logDuration instead of Duration to reduce the outlier effect in it; we attempt the same with Team and Advisors in case their relationship with the dependent variable is non-linear but find that the effect of outliers on the variables' relationship with logRaised is small, i.e. variables' pair-wise correlations with logRaised do not change by much when applying a log.

In Table 4.3 we examine the correlations between all continuous and other non-binary explanatory variables; we exclude from this analysis all pairs that include binary variables, as the correlation coefficient does not have a meaningful interpreta-

⁹By 'other non-binary variables' we refer to any of the following: integer variables such as Team and Advisors; discrete variables such as logDuration which by construction can only take the values log1, log2, ... and Rating which ranges between 1-5 and takes values up to the first or second decimal digit; variables that are theoretically continuous but exhibit significant clustering around certain numbers, such as Distribution.

Table 4.2: Token offerings sample statistics 2017 – 2020

	Mean	St. dev.	Skew.	Kurt.	Min.	Max.
logRaised	14.81	1.99	-0.98	4.78	5.24	20.72
logtCap	17.65	1.78	-0.34	16.28	0.65	33.75
Rating	3.4	0.6	-0.4	2.5	1.3	4.8
Team	9.54	6.54	1.92	9.14	1	54
Advisors	5.33	5.77	1.95	10.62	0	51
ETH price	420.82	244.01	0.73	2.85	10.10	1359.48
ETH mom	0.04	0.91	0.72	3.64	-2.65	3.98
Distribution	0.51	0.23	-0.36	2.61	0.00	1.00
logDuration	3.48	1.33	-1.09	4.2	0.00	6.78

Sample statistics for continuous and other non-binary variables based on a sample of 1,319 offerings completed between January 2017 - December 2020.

Table 4.3: Token offerings correlation matrix 2017 – 2020

	logRaised	logtCap	Rating	Team	Advisors	ETH price	ETH mom	Distribution
logtCap	0.36*	1						
Rating	0.17*	0.05	1					
Team	0.19*	0.06	0.20*	1				
Advisors	0.16*	0.06	0.31*	0.15*	1			
ETH price	0.21*	0.10*	-0.10*	-0.00	0.09*	1		
ETH mom	0.01	0.04	-0.10*	-0.04	-0.09*	0.28*	1	
Distribution	-0.04	-0.21*	-0.06	-0.03	0.02	0.10*	-0.02	1
logDuration	-0.08*	0.01	0.02	-0.01	0.05	-0.08*	-0.20*	0.37*

* $p < 0.01$

Correlation matrix of continuous and other non-binary variables based on a sample of 1,319 offerings completed between January 2017 - December 2020.

tion for such variables. As shown in the first column of Table 4.3, all variables exhibit significant correlations with logRaised, except for ETH mom and Distribution. Some of the pair-wise correlations between explanatory variables are significant, with the largest in magnitude being 0.37 between logDuration and Distribution. To ensure that multicollinearity is not an issue, we calculate the variance inflation factors (VIF) as defined e.g. in Hair (2010), that measure how much the variance of each estimated regression coefficient is increased because of collinearity with other explanatory variables. We find an average VIF across all variables of 1.26 and all VIFs are below 2, which is suggested as the strictest threshold value for identifying multicollinearity.

We also include similar analysis for the smaller sample of 902 capped offerings. Table 4.4 shows the sample statistics of continuous and other non-binary variables.

Table 4.4: Token offerings (capped) sample statistics 2017 – 2020

	Mean	St. dev.	Skew.	Kurt.	Min.	Max.
logSoftcap	14.49	1.49	-0.44	4.91	7.53	21.74
logtCap	17.57	1.88	-0.7	19.62	0.65	33.75
Rating	3.4	0.6	-0.5	2.6	1.3	4.8
Team	9.73	6.67	1.79	8.42	1	54
Advisors	5.44	5.83	1.92	10.39	0	51
ETH price	416.27	244.08	0.69	2.8	95.8	1359.48
ETH mom	-0.04	0.86	0.82	3.78	-2.33	3.29
Distribution	0.53	0.22	-0.41	2.82	0.00	1.00
logDuration	3.7	1.19	-1.23	5.25	0.00	6.78

Sample statistics of continuous and other non-binary variables based on a sample of 902 capped offerings completed between January 2017 - December 2020.

Table 4.5: Token offerings (capped) correlation matrix 2017 – 2020

	logSoftcap	logtCap	Rating	Team	Advisors	ETH price	ETH mom	Distribution
logtCap	0.41*	1						
Rating	0.04	0.08	1					
Team	0.16*	0.06	0.19*	1				
Advisors	0.10*	0.07	0.35*	0.17*	1			
ETH price	0.07	0.06	-0.13*	-0.00	0.05	1		
ETH mom	0.01	0.04	-0.10*	-0.00	-0.08	0.34*	1	
Distribution	-0.00	-0.24*	-0.07	-0.04	-0.01	0.07	-0.02	1
logDuration	0.03	0.01	-0.01	-0.07	-0.03	-0.17*	-0.21*	0.25*

* $p < 0.01$

Correlation matrix of continuous and other non-binary variables, based on a sample of 902 capped ICOs completed between January 2017 - December 2020.

As previously mentioned, log transformations are again applied to Softcap, tCap and Duration; the kurtosis of logSoftcap has a more reasonable value, as opposed to well above 100 before applying a log. Table 4.5 shows the pair-wise correlations between explanatory variables for the sample of 902 capped offerings; we do not include correlations with the dependent variable of soft cap exceedance as it is binary. Certain pair-wise correlations are significant but the VIFs of all variables are again below 2, suggesting that multicollinearity is not an issue.

Sample Period 2021 – 2022

Similarly, we show the sample statistics and correlation analysis for variables used in the 2021 – early 2022 sample in Tables 4.6 and 4.7. As before, we note some cases of high correlation between explanatory variables, but again the average and maximum VIF values are below the threshold value of 2 for identifying multicollinearity.

Table 4.6: Token offerings sample statistics 2021 – 2022

	Mean	St. Dev.	Skewness	Kurtosis	Min.	Max.
logRaised	12.03	1.23	1.72	7.75	8.01	18.51
logtCap	14.31	1.53	0.5	3.84	7.27	21.53
ETH price	3350.57	905.3	-0.42	2.12	822.32	4810.97
ETH mom	0	0.01	0.34	2.87	-0.02	0.03
Google	33.56	9.76	1.62	7.14	17.28	97.33
Offerings	2.78	1.93	3.13	18.25	1	15
Distribution-H	0.26	0.4	0.93	1.91	0	0.99
logDuration	0.2	0.56	3.14	13.25	0	4.25

Sample statistics for continuous and other non-binary variables based on a sample of 1,607 offerings completed between January 2021 - January 2022.

Table 4.7: Token offerings correlation matrix 2021 – 2022

	logRaised	logtCap	ETH price	ETH mom	Google	Offerings	Distribution-H
logtCap	0.38*						
ETH price	-0.00	-0.25*					
ETH mom	-0.01	0.19*	0.02				
Google	0.04	0.22*	-0.26*	0.35*			
Offerings	-0.11*	0.09*	0.19*	-0.08*	-0.04		
Distribution-H	-0.05	-0.06	0.03	-0.01	0.01	-0.04	
logDuration	0.14*	0.03	0.01	-0.02	-0.02	-0.03	0.00

* $p < 0.01$

Correlation matrix of continuous and other non-binary variables based on a sample of 1,607 offerings completed between January 2021 - January 2022.

4.4 Models

This Section first presents the model specifications used to obtain empirical results. We then formalise the interpretation of linear regression coefficients for certain special cases and also provide details on the standardisation process used for some of the

variables. Note that, as mentioned previously, we use separate model specifications for the 2017 – 2020 and 2021 – 2022 sample periods.

For the 2017 – 2020 sample period, we choose a linear regression model to connect the log amount raised for each token offering i with its possible determinants, as shown in equation (4.1). Since soft cap exceedance is a binary variable, we must use a binary choice model to connect it with possible determinants. Equation (4.2) shows the probit model we use.¹⁰

$$\begin{aligned} \log \text{Raised}_i = & \beta_0 + \beta_1 \log \text{Cap}_i + \beta_2 \text{Presale}_i + \beta_3 \text{Rating}_i + \beta_4 \text{Team}_i + \beta_5 \text{Advisors}_i + \\ & \beta_6 \text{ETH price}_i + \beta_7 \text{ETH mom}_i + \beta_8 \text{OFC}_i + \beta_9 \text{Bonus}_i + \beta_{10} \text{Distribution}_i + \\ & \beta_{11} \log \text{Duration}_i + \text{Controls}_i + \varepsilon_i \end{aligned} \quad (4.1)$$

$$\begin{aligned} \Pr[\text{Raised}_i \geq \text{Softcap}_i] = & \Phi(\gamma_0 + \gamma_1 \log \text{Softcap}_i + \gamma_2 \log \text{Cap}_i + \gamma_3 \text{Presale}_i + \gamma_4 \text{Rating}_i + \\ & \gamma_5 \text{Team}_i + \gamma_6 \text{Advisors}_i + \gamma_7 \text{ETH price}_i + \gamma_8 \text{ETH mom}_i + \\ & \gamma_9 \text{OFC}_i + \gamma_{10} \text{Bonus}_i + \gamma_{11} \text{Distribution}_i + \\ & \gamma_{12} \log \text{Duration}_i + \text{Controls}_i + \varepsilon_i) \end{aligned} \quad (4.2)$$

The variable names in equations (4.1) and (4.1) are as defined previously in sections 4.3 and 4.2, $\Pr$ denotes a probability and $\Phi(z)$ is the cumulative distribution function of the standard normal distribution. The terms β_1 , γ_1 and γ_2 are used to test $H1$ with an expected positive sign, except for γ_2 which is expected to have a negative sign; similarly, β_2 to β_5 and γ_3 to γ_6 are used for $H2$, with positive signs expected for genuine signals of public credibility and insignificant coefficients or even negative for cheap signals; β_6 , β_7 , γ_7 and γ_8 are used for $H3$ with an expected

¹⁰Note that we exclude the variables Github, Twitter and Whitepaper from both the linear and probit models because they produce inconsistent results.

positive sign; β_8 and γ_9 are used for $H4$ with expected positive signs; and β_9 to β_{11} and γ_{10} to γ_{12} are used for $H5$, for which it is expected that the coefficients will be insignificant. Control variables are as defined previously in Section 4.3.

Similarly, for the 2021 – 2022 sample period, we use a linear regression model as follows:

$$\begin{aligned} \log Raised_i = & \beta_0 + \beta_1 \log tCap_i + \beta_2 ETH\ price_i + \beta_3 ETH\ mom_i + \beta_4 Google_i + \\ & \beta_5 Offerings_i + \beta_6 Distribution-H_i + \beta_7 \log Duration_i + \\ & Type\ Controls_i + Category\ Controls_i + Platform\ Controls_i + \varepsilon_i \end{aligned} \quad (4.3)$$

Again, the variable names in equation (4.3) are as defined previously. The term β_1 is used to test $H1$ with an expected positive sign; β_2 to β_4 are used for $H3$ with expected positive signs; β_6 and β_7 are used for $H5$ and the Platform control variables are used for $H2$ as described in Section 4.2.¹¹

We now formalise the interpretation of linear regression coefficients when variables are used in log form. We do so in order to provide analysis on the magnitude of the variables' effects, as well as their sign and statistical significance; most papers in the relevant literature do not interpret the magnitude of regression coefficients, with the exceptions of Aggarwal et al. (2019), Albrecht et al. (2020), Lyandres et al. (2022), Blaseg (2018) and Lee et al. (2018) that provide brief coefficient interpretations but do not explain the impact of standardization or the application of a log.

Let us assume that a dependent variable $\log Y$ is used with a continuous independent variable $\log X$ in a linear regression model $\log Y = \alpha + \beta \log X + \varepsilon$. If the

¹¹In order to render the results comparable across different variables, we apply standardization within each sample period examined: 2017 – 2020 and 2021 – 2022. We standardize the dependent variable $\log Raised$ and the explanatory variables $\log tCap$, $\log Softcap$, $Distribution$, $ETH\ price$ and $ETH\ mom$, i.e. we de-mean them and impose a unit standard deviation. We do so because now the changes in these variables can be expressed in multiples of each variable's standard deviation and the estimated coefficients are comparable across these variables. We avoid standardizing binary variables and also variables with units that are already difficult to handle such as: Rating which is measured in a Likert-type scale; Team and Advisors which are measured in 'number of people'; $\log Duration$ which is measured in $\log(Days)$.

independent variable X changes by a factor δ , i.e. $X^* = \delta X$, then this translates to a change in Y by a factor δ^β , i.e.

$$Y^* = \delta^\beta Y. \quad (4.4)$$

This is because $\log Y^* = \alpha + \beta \log X^* + \varepsilon$ and $\log X^* - \log X = \log \delta$, so $\log Y^* - \log Y = \beta(\log X^* - \log X)$ and therefore $\log(\frac{Y^*}{Y}) = \beta \log \delta$.

However, we may have the same dependent variable in a linear regression model but with the independent variable not in log form, i.e. $\log Y = \alpha + \beta X + \varepsilon$. Then it is more reasonable to examine a change in X by δ units, i.e. $X^* = X + \delta$. This change translates to a change in Y by a factor $e^{\beta\delta}$, i.e.

$$Y^* = e^{\beta\delta} Y. \quad (4.5)$$

This is because $\log Y^* = \alpha + \beta X^* + \varepsilon$ and $X^* - X = \delta$, so $\log Y^* - \log Y = \beta(X^* - X)$ and therefore $\log(\frac{Y^*}{Y}) = \beta\delta$. Note that in this case, we often examine a change in X by $\delta = 1$ unit which translates to a change in Y by a factor of e^β . Equations (4.4) and (4.5) hold for multivariate linear regressions, provided that all other independent variables are unchanged.

We do not encounter such issues in probit regressions, as we calculate the variables' marginal effects explicitly due to the non-linearity of the probit model. Specifically, we calculate the marginal effect at means (MEM), i.e. the effect on soft cap exceedance probability from a one unit change in a non-binary independent variable or a transition from 0 to 1 for a binary independent variable, when all other independent variables are fixed at their mean values.

4.5 Empirical Results

In the following, we discuss our findings in relation to the relevant literature. We first provide and discuss the results for the January 2017 – December 2020 period, followed by the more recent January 2021 – January 2022 period.

4.5.1 Sample Period 2017 – 2020

The linear regression model results with logRaised as the dependent variable are shown in Table 4.8. We use the sample of 1,319 offerings as described in Section 4.3, which excludes approximately 200 offerings with zero funding raised, to avoid bias in the model’s estimates. We present results for three sub-periods: 1 January 2017 – 30 June 2018, containing 589 offerings; 1 July 2018 – 31 December 2019, containing 626 offerings; 1 January 2020 – 31 December 2020, containing 104 offerings; for completeness, the entire 2017 – 2020 sample results are shown in Table C1 of Appendix C. We expect the results on the first sub-period to be consistent with the findings in the relevant literature due to significant overlap in the sample period.

We attempt a basic model configuration (i) which only includes the explanatory variables related to our hypotheses and a full configuration (ii) which further includes control variables; specifically for the third sub-period January – December 2020 which only includes a small sample of 104 offerings, we only attempt the basic configuration to avoid overfitting. Following common practice in the relevant literature we examine three levels of significance at 10%, 5% and 1%. The regressions are estimated in Stata using the ordinary least squares (OLS) method with heteroscedasticity-robust standard errors; the standard errors are estimated using the Huber/White/sandwich estimator, as described e.g. by Cameron and Trivedi (2009).

The linear model results are consistent across both configurations with almost all variables that are significant in the basic model remaining significant in the full model and the adjusted R^2 very slightly improving in the full model. In the full model setup using the first sub-period of the sample, the adjusted R^2 of 0.286 is

mostly consistent with the results of the existing literature, considering that the majority of relevant papers end their sample period in Q2, Q3 or Q4 of 2018, e.g. with the sample of Lyandres et al. (2022) going up to November 2018 as shown in Table 4.1. Interestingly, the adjusted R^2 is reduced significantly to 0.20 in the second sub-period and even more so to 0.12 in the third sub-period, suggesting that in the post-boom and Covid-19 pandemic periods, token offering characteristics are less suitable as success determinants. Specifically for the third sub-period, we observe a marked decrease between the R^2 and its adjusted value from 0.21 to 0.11, indicating no more explanatory variables should be used due to the small sample size. Note that in all cases, the corresponding F-tests based on the estimated R^2 values indicate that the models presented in Table 4.8 outperform the constant-only model.

The probit regression model results with probability of soft cap exceedance as the dependent variable are shown in Table 4.9, where we report the marginal effects at means (MEM), i.e. the effect on soft cap exceedance probability from a one unit change in each independent variable when all other independent variables are fixed at their mean values. We use the smaller sample of 902 capped offerings, as described in Section 4.3. This sample is again split in three sub-periods, with the entire sample results shown in Table C2 of Appendix C. Note that as a result of showing the marginal effects instead of the outright estimated coefficients, the results in Table 4.9 do not include the constant coefficient.

The probit regressions are estimated with the maximum likelihood estimation (MLE) method and without the use of robust errors, as suggested by Cameron and Trivedi (2009) to avoid the danger of misspecification in binary choice models. We follow again the basic and full model setups considering any variable with a p-value above 10% to be insignificant and find that results are mostly consistent across setups, with some differences in the level of significance. Again we only estimate the basic model specification for the third sub-period to avoid overfitting.

In evaluating the goodness-of-fit for the probit regression model, we use McFadden's pseudo- R^2 as discussed e.g. by Cameron and Trivedi (2009). It is defined as

Table 4.8: Linear regression model parameters 2017 – 2020

	Jan 2017 - Jun 2018				Jul 2018 - Dec 2019				Jan - Dec 2020	
	(i) Basic		(ii) Full		(i) Basic		(ii) Full		(i) Basic	
Constant	-0.931***	(-3.83)	-0.937***	(-3.38)	-0.224	(-1.00)	0.098	(0.35)	-1.074*	(-1.81)
logtCap	0.254***	(4.59)	0.264***	(5.06)	0.360***	(8.70)	0.348***	(8.01)	0.317**	(2.58)
Presale	-0.036	(-0.49)	-0.042	(-0.57)	-0.021	(-0.28)	-0.012	(-0.15)	-0.249	(-1.21)
Rating	0.393***	(6.11)	0.354***	(5.45)	0.088	(1.39)	0.100	(1.52)	0.100	(0.85)
Team	0.022***	(4.14)	0.021***	(3.87)	0.017***	(3.14)	0.018***	(3.10)	0.024**	(2.26)
Advisors	0.005	(0.83)	0.004	(0.71)	0.007	(1.08)	0.009	(1.22)	-0.003	(-0.11)
ETH price	0.167***	(3.83)	0.165***	(3.75)	0.166***	(2.84)	0.162***	(2.70)	-0.266*	(-1.77)
ETH mom	-0.010	(-0.32)	0.000	(0.01)	-0.048	(-0.86)	-0.027	(-0.45)	0.175	(1.64)
OFC	0.162**	(2.23)	0.166**	(2.27)	0.211***	(2.97)	0.193***	(2.59)	0.064	(0.33)
Bonus	-0.131*	(-1.80)	-0.150**	(-2.03)	-0.108	(-1.38)	-0.109	(-1.35)	-0.296	(-1.26)
Distribution	-0.081*	(-1.74)	-0.079	(-1.61)	0.113**	(2.54)	0.092**	(2.03)	-0.009	(-0.06)
logDuration	-0.149***	(-4.65)	-0.165***	(-4.95)	-0.071***	(-2.62)	-0.085***	(-2.84)	0.019	(0.32)
Accepts BTC			0.149*	(1.79)			0.157*	(1.95)		
Accepts ETH			-0.018	(-0.12)			-0.178	(-1.02)		
Accepts other			0.002	(0.02)			-0.022	(-0.25)		
Accepts fiat			0.192*	(1.80)			-0.162	(-1.57)		
Capped			0.171	(1.35)			-0.210	(-1.39)		
Ethereum-based			0.033	(0.33)			0.006	(0.05)		
IEO			0.386	(1.01)			-0.204*	(-1.69)		
KYC			-0.027	(-0.29)			0.046	(0.50)		
Whitelist			0.119	(1.33)			-0.009	(-0.13)		
High			0.006	(0.07)			0.092	(1.22)		
Medium			-0.156**	(-2.08)			-0.075	(-1.05)		
Low			0.043	(0.43)			-0.073	(-0.64)		
Observations	589		589		626		626		104	
R^2	0.292		0.314		0.209		0.226		0.210	
Adj. R^2	0.278		0.286		0.195		0.197		0.116	

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Estimated parameters for the linear regression model with logRaised as the dependent variable. The sample period is divided into three sub-periods: 1 January 2017 - 30 June 2018, 1 July 2018 - 31 December 2019 and 1 January - 31 December 2020. The model specifications are: (i) Basic which only includes the variables involved in our hypotheses; (ii) Full which further includes control variables. The first column in each model specification displays the regression betas and the second column displays the t-statistics in parentheses.

$R^2_{\text{McFadden}} = 1 - \frac{\log L}{\log L_c}$, where $\log L$ is the maximum log likelihood of our probit model and $\log L_c$ is the value of the log likelihood in the intercept-only probit model. The estimated value of McFadden's pseudo- R^2 is consistent across both model setups, again improving in the full model. The full model specification's pseudo- R^2 in the first sub-period of the sample is 0.193, somewhat consistent with the findings in the relevant literature, considering the differences in sample size: for instance, Roosenboom et al. (2020) use a sample of 185 offerings until December 2017 in a logit regression with the same dependent variable and find a pseudo- R^2 of 0.233; Bourveau et al. (2022) use a similar model setup with a sample of 341 offerings until February 2018 and find a pseudo- R^2 of 0.31. When examining the second sub-period, the

Table 4.9: Probit regression marginal effects 2017 – 2020

	Jan 2017 - Jun 2018				Jul 2018 - Dec 2019				Jan - Dec 2020	
	(i) Basic		(ii) Full		(i) Basic		(ii) Full		(i) Basic	
logSoftcap	-0.136***	(-4.09)	-0.144***	(-4.22)	-0.138***	(-4.83)	-0.147***	(-4.99)	-0.148	(-1.59)
logtCap	0.058*	(1.72)	0.057*	(1.69)	0.095***	(3.23)	0.101***	(3.23)	0.048	(0.60)
Presale (d)	-0.027	(-0.46)	-0.024	(-0.39)	-0.025	(-0.52)	-0.031	(-0.64)	-0.101	(-0.74)
Rating	0.264***	(4.86)	0.255***	(4.50)	0.020	(0.51)	0.036	(0.85)	-0.021	(-0.19)
Team	0.005	(1.21)	0.005	(1.20)	0.008**	(2.06)	0.009**	(2.22)	0.002	(0.18)
Advisors	0.001	(0.11)	0.000	(0.06)	0.004	(0.92)	0.006	(1.25)	0.006	(0.24)
ETH price	0.052	(1.31)	0.034	(0.82)	0.035	(0.83)	0.035	(0.78)	0.152	(0.78)
ETH mom	0.006	(0.26)	0.012	(0.46)	-0.015	(-0.44)	0.014	(0.36)	0.036	(0.36)
OFC (d)	0.137**	(2.27)	0.155**	(2.50)	0.056	(1.18)	0.062	(1.28)	0.004	(0.02)
Bonus (d)	-0.117**	(-2.06)	-0.123**	(-2.10)	-0.033	(-0.69)	-0.026	(-0.51)	-0.189	(-1.21)
Distribution	-0.036	(-1.08)	-0.036	(-1.03)	0.004	(0.16)	-0.000	(-0.01)	-0.008	(-0.11)
logDuration	-0.078**	(-2.23)	-0.097***	(-2.65)	-0.026	(-1.18)	-0.047*	(-1.94)	-0.027	(-0.66)
Accepts BTC (d)			0.115*	(1.67)			0.105*	(1.80)		
Accepts ETH (d)			-0.153	(-1.47)			-0.011	(-0.10)		
Accepts other (d)			0.051	(0.72)			-0.073	(-1.21)		
Accepts fiat (d)			0.083	(0.94)			0.012	(0.20)		
Ethereum-based (d)			0.008	(0.11)			0.008	(0.10)		
IEO (d)							-0.204**	(-2.15)		
KYC (d)			0.095	(1.33)			-0.054	(-0.94)		
Whitelist (d)			-0.051	(-0.70)			-0.021	(-0.40)		
High (d)			-0.070	(-1.19)			0.021	(0.42)		
Medium (d)			-0.064	(-1.07)			-0.048	(-1.00)		
Low (d)			0.087	(1.09)			-0.029	(-0.42)		
Observations	366		366		474		474		62	
Pseudo- R^2	0.161		0.193		0.055		0.072		0.215	

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Marginal effects at means (MEM) for the probit regression with soft cap exceedance as the dependent variable. The sample period is divided into three subperiods: 1 January 2017 - 30 June 2018, 1 July 2018 - 31 December 2019, 1 January 2020 – 31 December 2020. The model specifications are: (i) Basic which only includes the variables involved in our hypotheses; (ii) Full which further includes control variables. The notation (d) next to a variable indicates that the corresponding marginal effect is calculated for a discrete change of the binary variable from 0 to 1 and the constant term is not included as MEM is not defined for constants. The first column in each model specification displays the marginal effect and the second column displays the z-statistics in parentheses.

pseudo- R^2 is reduced by more than half to 0.072, again suggesting that token offerings characteristics are much less suitable for use as success determinants. Surprisingly, in the third sub-period the pseudo- R^2 is 0.22 but the null hypothesis that the model specification outperforms a constant-only model cannot be rejected even at a 10% significance level based on a likelihood ratio test statistic, so we do not consider this value to be reliable; for comparison, we note that all model specifications in the first and second sub-periods reject the null hypothesis.

We now discuss the results of both models for each research hypothesis in turn and compare them to the findings of papers in the relevant literature. Note that we refer to the parameter estimates for the basic model (i) in tables 4.8 and 4.9, unless

otherwise specified. Also note that we interpret parameter estimates and marginal effects under the assumption that all other independent variables remain unchanged.

Consistent with our hypothesis (*H1*) that a project's target market cap is a good predictor of the actual amount raised in the token sale, the logtCap coefficient is significant and positive in all model specifications and all sub-periods; we further find that its effect is more pronounced during the post-boom and pandemic periods.

In the first sub-period of our sample, a 1-standard-deviation increase in logtCap causes, all else being equal, a change of 0.254 standard deviations in logRaised, since both the dependent variable and logtCap are standardized. According to equation (4.4), we interpret the coefficient as follows: a 1% increase in tCap causes an increase of 0.263% in the amount raised, since both variables are in log form. As mentioned in Section 4.2, we cannot find any source in the relevant literature that uses this variable, but some use the two components of logtCap, Price and Supply, in regressions that explain logRaised. Roosenboom et al. (2020) find that $\log(1 + Price)$ is insignificant in explaining $\log(1 + Raised)$ but they find the coefficient of $\log(1 + Supply)$ to be positive and significant; Momtaz (2020a), Fisch (2019) and Amsden and Schweizer (2018) also find the same for $\log(Supply)$ – see Table 4.1. In the second and third sub-periods of our sample, the amount raised is more sensitive to changes in the target market cap: a 1% increase in tCap causes the amount raised to increase respectively by 0.359% and 0.316%.

Similarly, logtCap is significant in explaining the soft cap exceedance probability in the probit regression; the effect is more pronounced during the post-boom period but completely vanishes during the Covid-19 pandemic. During the first sub-period, its marginal effect (at means) indicates that a 1-standard-deviation change in logtCap increases the probability of soft cap exceedance by 0.058, i.e. by 5.8 percentage points, assuming that all other explanatory variables are fixed at their mean values. This somewhat contradicts the finding of Roosenboom et al. (2020) that neither Price nor Supply are significant in explaining soft cap exceedance in a logit model; using a similar model setup, Rasskazova et al. (2019) find that Price has a negative effect

on soft cap exceedance, which partially confirms our result. When examining the second sub-period of the sample, a 1-standard-deviation change in logtCap increases the probability of soft cap exceedance by 9.5 percentage points.

In *H1* we further propose that a higher soft cap level has a negative effect on the probability of soft cap exceedance; indeed, we find a significant negative coefficient for logSoftcap and the effect is again more pronounced during the post-boom period. Its marginal effect as shown in Table 4.9 indicates that a 1-standard deviation increase in logSoftcap causes the probability of soft cap exceedance to decrease by 13.6 percentage points in the first half of the sample and by 13.8 percentage points in the second half. Roosenboom et al. (2020) also finds that a higher soft cap has a negative effect on the probability of soft cap exceedance, confirming our result.

Given that both logtCap and logSoftcap are standardized and their marginal effects coefficients are directly comparable, it seems that the magnitude of the soft cap's effect on the probability of soft cap exceedance is larger than that of the target market cap, suggesting that setting a low soft cap is more important than accurately choosing the token's target market cap.

In *H2* we assume that the amount raised and the soft cap exceedance probability increase with signals of public credibility, as long as the signals are genuine. In Section 4.2 we identify Rating and Presale as genuine signals and argue that Team and Advisors can be considered as cheap signals.

Interestingly, Rating has a significant positive effect on logRaised and the probability of soft cap exceedance during the initial boom period of 2017 – mid-2018 but *not* in the subsequent periods. The insignificance of Rating after June 2018 indicates that investors may be taking notice of the allegations made by (Hartmann, 2018, published on Medium on 14 June 2018), the warnings of Boreiko and Vidusso (2019) that less trust should be placed on token offering ratings provided by ICObench and similar aggregator websites and also the warning of moral hazard in venture credibility signals by Momtaz (2020a).

In the first sub-period of our sample, the Rating coefficient in the linear model is positive and strongly significant; all else being equal, an increase in Rating by 1 unit causes an increase of 0.393 standard deviations in logRaised, i.e. the amount raised increases by 48.1% according to equation (4.5), since the Rating variable is not in log form.

To justify this apparently very strong effect, note that we do not standardize Rating and therefore a 1-unit change is quite significant in the 1-5 scale used in the ICObench rating. The findings of papers in the relevant literature on Rating mostly support our result for the linear model, with Boreiko and Vidusso (2019) also finding a positive significance using the exact same variable and Roosenboom et al. (2020), Aggarwal et al. (2019) and Fisch (2019) obtain similar results using individual components of the ICObench rating. Only Momtaz (2020a) finds the expert rating component to be insignificant and Fisch (2019) finds the project vision component to be negatively significant. Similarly, the effect of Rating on the probability of soft cap exceedance is significant only during the first sub-period of our sample. Assuming that all other variables are fixed at their mean values, an increase in Rating by 1 unit improves the probability of soft cap exceedance by 0.264, i.e. 26.4 percentage points. When examining the second and third sub-periods of our sample, the coefficient of Rating is not significant in explaining logRaised or the probability of soft cap exceedance, even at a 10% level.

The effect of holding a token presale prior to the public offering has no effect on logRaised or the probability of soft cap exceedance, both during and after the 2017 boom period, even at a 10% significance level. Given the mixed results in the literature regarding Presale, our finding is in agreement with certain authors, e.g. with Fisch, 2019, and in disagreement with others. Contrary to our finding, Roosenboom et al. (2020) and Rasskazova et al. (2019) both find that Presale is significant in explaining soft cap exceedance probability but disagree on the sign of its coefficient.

The variable Team maintains its significance during and after the 2017 boom period and even during the Covid-19 pandemic in the linear model; in the probit model its coefficient is insignificant in the first sub-period of the sample but significant and positive in the second sub-period. In the first sub-period of the sample, if the size of the venture's team increases by 1, i.e. one more member is added to the team, logRaised increases by 0.022 standard deviations and therefore the amount raised increases by approximately 2.22%, according to equation (4.5). There is overwhelming consensus over the positive significance of Team in the linear model, as confirmed by Roosenboom et al. (2020), Amsden and Schweizer (2018), Aggarwal et al. (2019), Ante and Fiedler (2020), Chen (2019), Lyandres et al. (2022) and Bourveau et al. (2022). Bourveau et al. (2022) confirm our finding that team size is insignificant in explaining the probability of soft cap exceedance, while Roosenboom et al. (2020) and Rasskazova et al. (2019) find it to be significant. When examining the second sub-period of our sample, an increase in team size by one member causes a slightly smaller increase of 1.75% in the amount raised and an increase of 0.8 percentage points in the probability of soft cap exceedance; during the Covid-19 pandemic period, a one-member increase in team size causes a 2.43% increase in the amount raised. The effect of Advisors is insignificant in all sub-periods, in both the linear and the probit model. The results in the relevant literature are mixed but mostly support our finding: Amsden and Schweizer (2018) find the effect of advisory team size to be positive and significant in one linear model setup but insignificant in another and Ante and Fiedler (2020) also find no significance.

We reiterate the importance of our finding that the token offering rating has no explanatory power in the post-boom and pandemic periods, suggesting that investors may be reducing their trust in ratings provided by ICObench and similar aggregator websites. In the context of distinguishing between genuine and cheap credibility signals, our results are mixed: Rating, a seemingly genuine signal, is significant during the 2017 boom period but not so in the post-boom and pandemic periods; Presale, another genuine signal, is always insignificant; venture team size,

a cheap signal is mostly significant; advisory team size, another cheap signal, is always insignificant. These results confirm the findings of Ante and Fiedler (2020), suggesting that investors consider certain cheap signals to be reliable and may not trust other signals, even if they are costly to produce.

The ETH price variable, i.e. the average ETH/USD price during each offering, has a strong positive significance in the linear model in during and after the 2017 boom period, indicating that it captures investor trust in the entire crypto asset class; however, its coefficient is significant and negative during the pandemic, suggesting that offerings raised more funding in 2020 when ether's price was lower. A 1-standard-deviation increase in ETH price increases the amount raised by 18.2% during the 2017 boom period and by 18.1% in the post-boom period according to equation (4.5); however during the pandemic period the same increase would cause the amount raised to decrease by 23.4%. Moreover, ETH price is insignificant in explaining the probability of soft cap exceedance and the variable ETH mom, i.e. the ETH/USD average momentum, is insignificant in all model specifications. Amsden and Schweizer (2018) find mixed results for ETH price, i.e. positively significant in one linear model setup and insignificant in another and Boreiko and Vidusso (2019) confirm our finding that ETH mom has no explanatory power.

Our hypothesis on investor preference for tax-haven-domiciled ventures is partially confirmed; the OFC (offshore financial centre) variable remains significant in the linear model before and after the 2017 boom but its effect is insignificant during the pandemic period; in the probit model, it is only significant during the 2017 boom. If a venture is domiciled in a tax-haven, the amount raised increases by 17.6% in the first sub-period of the sample and by 23.5% in the second sub-period, according to equation (4.5). Similarly, in the first sub-period of the sample the soft cap exceedance probability increases by 0.137, i.e. 13.7 percentage points. This result is in contradiction with all relevant findings in the literature, as Amsden and Schweizer (2018) and Ante and Fiedler (2020) both find the tax-haven domicile variable to be insignificant in explaining ICO success and Huang et al. (2020) find that tax-havens

do not attract more token offerings compared with other jurisdictions. To justify this contradiction, we note that such results very much hinge on the definition of tax havens; as explained in Section 4.3, we choose to include Singapore as a tax haven and this almost doubles the population of tax-haven-domiciled ventures in our sample.

Our findings on the coefficients of the Bonus, Distribution and logDuration variables show that they are important in the 2017 boom period, less important in the post-boom period and insignificant during the pandemic period, partially confirming our hypothesis that they have no effect on the amount raised or the probability of soft cap exceedance. Interestingly, the sign of Distribution's coefficient flips from negative in the first sub-period to positive in the second.

In the first sub-period of the sample, all three variables have significant and negative coefficients in the linear model: the inclusion of a bonus scheme decreases the amount raised by 12.3%; an increase in Distribution by 1 standard deviation decreases the amount raised decreases by 7.7%; an increase in logDuration by 1 unit also causes a decrease in the amount raised by 13.8%. When examining the probit model results for the first sub-period of the sample, only Bonus and logDuration are significant: the inclusion of a bonus scheme decreases the soft cap exceedance probability by 11.6 percentage points; a 1-unit increase in logDuration decreases the soft cap exceedance probability by 7.8 percentage points.

We note that it is difficult to interpret what a change of 1 unit in logDuration means for the original variable Duration. An illustrative example is that when Duration increases by 1 day from its average value of 59 days, this translates to an increase in logDuration of approximately 0.02 units; so based on the 2017 boom period results mentioned, for the average token offering that lasts approximately 2 months, a 1-day increase in duration can decrease the amount raised by almost 30%.

The results in the relevant literature are mixed for Bonus and Distribution but quite similar for Duration. Roosenboom et al. (2020) find a significant coefficient for Bonus in their linear model but find the variable to be insignificant in their logit

model, while Amsden and Schweizer (2018) also find an insignificant coefficient for Bonus in a linear model. Amsden and Schweizer (2018) and Lyandres et al. (2022) find Distribution to be significant, but Fisch (2019) finds the opposite. Roosenboom et al. (2020) also find that $(1 - \text{Distribution})$, i.e. insider token retention, is insignificant. The consensus on the negative significance of offering duration is quite strong across papers in the relevant literature, see e.g. the discussions of Momtaz (2020a), Roosenboom et al. (2020), Chen (2019), Fisch (2019) and Rasskazova et al. (2019) in Section 4.2. When examining the second sub-period, Bonus is insignificant, partially confirming our hypothesis. Distribution and logDuration have significant coefficients in the second sub-period but only in the linear model. Interestingly, a 1-standard-deviation increase in Distribution causes the amount raised to increase by 12%, suggesting that when ventures offer a larger share of the token’s total supply to investors, they tend to raise more funding. A 1-unit increase in logDuration in the second sub-period causes a 6.9% decrease in the amount raised.

Finally, we present a discussion of the control variables. Note that we do not include control variables in the third sub-period to avoid overfitting. Accepts BTC, Accepts fiat, IEO and Medium are the only ones with significant coefficients.

Accepts BTC has a positive coefficient in both models and across both the first and second sub-periods. Its effect is more pronounced in the post-boom period for the linear model but less pronounced for the probit model; this is somewhat surprising given that tokens in our sample are almost exclusively based on the Ethereum platform; perhaps investors are negatively predisposed towards a token offering if it does not accept bitcoin as a means of payment, since bitcoin is the most influential crypto asset in the entire asset class. For instance, Rhue (2018) also finds that Accepts BTC is positively significant but finds the same for Accepts ETH, which we do not.

As previously mentioned, nearly all of the tokens in our sample are Ethereum-based which perhaps explains why Accepts ETH and also Ethereum-based are not significant variables, so investors possibly take these characteristics as a given. On the

other hand, this suggests that tokens deployed on other smart contract platforms such as EOS or NEO, are not at a disadvantage by the mere fact that they are not deployed on Ethereum. Moreover, there are mixed findings about the Ethereum-based variable in the relevant literature: Roosenboom et al. (2020) find it to be insignificant in a logit model setup but positively significant in a linear model; Amsden and Schweizer (2018) finds it to be insignificant in one linear setup and negatively significant in another; Momtaz (2020*a*), Boreiko and Vidusso (2019) and Fisch (2019) find a significant positive coefficient.

The effect of IEO is significant and negative but only in the post-boom period between mid-2018 – 2019.¹² Out of the industry-based binary variables High, Medium and Low, only Medium is significant in the linear model and only during the 2017 boom period, indicating that the venture’s self-declared business categories have no effect at all in the post-boom period. This suggests ventures more likely to be invested in by medium-income investors such as retail investors and households raised fewer funds during the boom period but this does not impact the probability of soft cap exceedance; in the post-boom period, the project categories appear to have no effect on fundraising success.

Finally, acceptance of fiat currency contributions is only positively significant during the 2017 boom and only in the linear model, as also found by Momtaz (2020*b*), Roosenboom et al. (2020) and Amsden and Schweizer (2018). We find that Capped is not a significant variable and the findings of papers in the relevant literature are again mixed. Our findings that KYC and Whitelist are insignificant agree with those of almost all papers in the literature, except for Lyandres et al. (2022) who find that both these variables are significant in explaining the amount raised.

¹²The IEO variable is excluded from the probit model in the 2017 boom period of the sample due to high collinearity.

4.5.2 Sample Period 2021 – 2022

The linear regression model results for the sample of 1,607 token offerings completed between January 2021 - January 2022 and with logRaised as the dependent variable are shown in Table 4.10. Given that this recent period of the market is not examined at all in the extant literature, we compare these results with our previous findings on the fundraising success determinants of token offerings from the 2017 - 2020 period. We attempt a basic model configuration (i) which only includes the main explanatory variables used in the 2017 - 2020 period, to the extent that they are available. We then introduce additional controls for: the type of token offering (ii), i.e. whether it is an IDO, an IEO or an ICO; the venture's category (iii), such as gaming, NFT, DeFi etc.; the launchpad platform (iv) used for each token offering, such as PancakeSwap, Polkastarter, Binance etc. Finally, we include all of the above variables and controls in the full model configuration (v). Again following common practice in the relevant literature we examine three levels of significance at 10%, 5% and 1% and estimate the linear regressions with OLS using heteroscedasticity-robust standard errors.

As shown in Table 4.10, the results are consistent across all configurations with almost all variables that are significant in the basic model remaining significant in the full model. Interestingly, the inclusion of the Type and Platform controls provides significant additional explanatory power, with an equally significant improvement in the adjusted R^2 , from 0.20 in the Basic configuration to 0.54 in the Full configuration. It is important to note that such a high adjusted R^2 is not exhibited in any of the results in the existing literature, nor in our own results for the 2017 - 2020 period of the ICO market. This suggests that in the most recent IDO boom period of 2021 - 2022, the platform used to conduct a token offering together with the type of offering prove to be most significant determinants of ICO fundraising success. Note also that the corresponding F-tests based on the estimated R^2 values indicate that the models presented in Table 4.10 outperform the constant-only model in all cases.

Table 4.10: Linear regression parameters for the 2021 – early 2022 period

	(i) Basic		(ii) Type		(iii) Category		(iv) Platform		(v) Full	
Constant	0.245***	(5.25)	1.979***	(9.69)	0.309***	(3.09)	0.797***	(7.63)	1.896***	(8.12)
logtCap	0.444***	(11.60)	0.401***	(10.97)	0.428***	(10.86)	0.365***	(10.99)	0.315***	(8.56)
ETH price	0.152***	(5.70)	0.179***	(7.01)	0.154***	(5.10)	0.125***	(4.86)	0.165***	(6.28)
ETH mom	-0.108***	(-4.41)	-0.139***	(-5.93)	-0.104***	(-4.03)	-0.102***	(-4.57)	-0.136***	(-6.17)
Google	0.013	(0.59)	0.026	(1.24)	0.015	(0.67)	0.016	(0.83)	0.018	(0.92)
Offerings	-0.094***	(-7.10)	-0.097***	(-8.08)	-0.096***	(-7.41)	-0.068***	(-6.23)	-0.062***	(-5.47)
Distribution-H	-0.095*	(-1.77)	-0.046	(-0.90)	-0.082	(-1.49)	-0.086*	(-1.86)	-0.044	(-0.97)
logDuration	0.206***	(4.53)	0.191***	(4.18)	0.214***	(4.56)	0.106**	(2.41)	0.109**	(2.55)
Observations	1607		1607		1607		1607		1607	
R ²	0.206		0.331		0.230		0.515		0.579	
Adj. R ²	0.202		0.327		0.208		0.488		0.544	
Type										
ICO (54)			0.000						0.000	
IDO (1340)			-1.731*** (-8.28)						-1.353*** (-5.74)	
IEO (213)			-2.200*** (-10.26)						-2.242*** (-4.84)	
Category										
Artificial (7)					0.712 (1.22)				0.324 (0.87)	
Blockchain (153)					0.000				0.000	
Business (3)					-0.162 (-1.14)				-0.132 (-0.92)	
Cloud (9)					0.199 (0.76)				0.082 (0.60)	
Crowdfunding (3)					-0.455** (-2.55)				-0.149 (-0.66)	
Dapp (14)					-0.251 (-0.69)				-0.195 (-0.92)	
Data (16)					0.554** (2.33)				0.122 (0.56)	
DeFi (411)					-0.079 (-0.76)				-0.039 (-0.46)	
Education (2)					-0.541** (-2.33)				0.224 (1.12)	
Energy (2)					0.809*** (7.43)				0.715*** (3.87)	
Entertainment (32)					0.352 (1.33)				-0.155 (-1.08)	
Exchange (57)					-0.045 (-0.33)				0.031 (0.29)	
Fan (6)					0.297 (0.56)				-0.846*** (-4.14)	
Financial (34)					-0.266* (-1.86)				-0.291** (-2.26)	
Gambling (13)					-0.432** (-2.03)				-0.220 (-1.11)	
Gaming (586)					-0.067 (-0.66)				-0.107 (-1.26)	
Healthcare (1)					0.269** (2.57)				0.139 (0.81)	
IOT (1)					-0.118 (-1.05)				-0.200 (-0.67)	
Insurance (10)					-0.140 (-0.55)				-0.215 (-1.26)	
Market (3)					-0.475** (-2.23)				-0.228** (-2.02)	
Marketplace (36)					-0.050 (-0.30)				0.040 (0.37)	
Media (8)					-0.234 (-0.93)				-0.024 (-0.13)	

	(i) Basic	(ii) Type	(iii) Category	(iv) Platform	(v) Full
Meme (2)			-0.512 (-0.73)		0.215 (0.57)
Mining (1)			-0.582*** (-6.29)		-0.242** (-2.50)
Network (5)			0.117 (0.15)		0.179 (0.44)
Non-Fungible (92)			-0.209* (-1.83)		-0.073 (-0.80)
Other			-0.720*** (-6.90)		-0.744*** (-3.27)
Predictions (3)			-0.395*** (-3.80)		-0.202 (-0.58)
Protocol (15)			0.142 (0.47)		0.094 (0.46)
Real (3)			-0.278 (-0.87)		-0.174 (-0.92)
Security (4)			-0.075 (-0.30)		0.016 (0.16)
Social (17)			0.462 (1.28)		0.306 (0.87)
Tokenized (1)			0.247** (2.49)		-0.645*** (-6.63)
Trading (28)			-0.274* (-1.65)		-0.015 (-0.13)
VR (15)			-0.466*** (-2.80)		-0.266** (-2.06)
Verification (1)			0.409*** (3.62)		0.372*** (3.93)
Wallet (11)			-0.283 (-0.56)		-0.609*** (-2.67)
Platform					
A2DAO (DEX, 15)				-1.003*** (-6.68)	-0.692*** (-4.78)
AcceleRaytor (DEX, 13)				-0.079 (-0.41)	0.230 (1.18)
AscendEX (CEX, 5)				-1.160** (-2.57)	0.092 (0.15)
AvaXlauncher (DEX, 4)				-0.701 (-1.61)	-0.489 (-1.12)
Avalaunch (DEX, 13)				0.368** (2.44)	0.635*** (4.02)
BSCPad (DEX, 57)				-0.553*** (-4.65)	-0.283** (-2.38)
BSCStation (DEX, 36)				-0.773*** (-7.46)	-0.575*** (-5.10)
BSClaunch (DEX, 13)				-0.907*** (-5.86)	-0.792*** (-4.04)
BinStarter (DEX, 12)				-0.509*** (-4.08)	-0.273** (-2.20)
Binance (CEX, 7)				1.486*** (7.71)	3.032*** (6.48)
BitMart (CEX, 1)				-2.111*** (-20.56)	-1.005** (-2.43)
Bounce (DEX, 14)				-1.613*** (-8.96)	-1.316*** (-7.40)
BullPerks (DEX, 13)				-0.838*** (-5.49)	-0.597*** (-3.67)
Bybit (CEX, 7)				-0.019 (-0.11)	1.181*** (2.66)
CardStarter (DEX, 12)				-0.279* (-1.96)	0.100 (0.65)
ChainBoost (DEX, 18)				-0.907*** (-6.57)	-0.602*** (-4.24)
Conv. Fin. (DEX, 1)				-1.722*** (-12.63)	-1.399*** (-9.73)
CyberFi (DEX, 21)				-0.633*** (-3.18)	-0.350* (-1.76)
DAO Maker (DEX, 80)				-0.608*** (-5.15)	-0.289** (-2.41)
Dodo (DEX, 2)				-0.286 (-0.55)	-0.189 (-0.41)
DuckSTARTER (DEX, 50)				-0.875*** (-7.89)	-0.574*** (-5.12)
ETHPad (DEX, 6)				-0.466** (-2.24)	-0.263 (-1.21)
Enjinstarter (DEX, 26)				-0.825*** (-7.89)	-0.575*** (-5.55)
FTX (CEX, 6)				-0.066 (-0.33)	1.244*** (2.81)
FantomStarter (DEX, 5)				-0.910*** (-5.68)	-0.565*** (-3.37)
FireStarter (DEX, 3)				-0.265 (-1.51)	-0.011 (-0.05)

	(i) Basic	(ii) Type	(iii) Category	(iv) Platform	(v) Full
GameFi (DEX, 35)				-0.951*** (-6.54)	-0.670*** (-4.52)
GameStation (DEX, 17)				-0.816*** (-5.21)	-0.579*** (-3.85)
GameZone (DEX, 10)				-0.187 (-1.03)	0.045 (0.24)
Gamestarter (DEX, 13)				-0.573*** (-2.84)	-0.321 (-1.61)
Gate.io (CEX, 134)				-1.493*** (-12.55)	-0.300 (-0.73)
GenPad (DEX, 3)				-0.852*** (-4.28)	-0.515*** (-2.67)
Gnosis (DEX, 2)				1.502*** (6.43)	1.718*** (7.79)
Huobi (CEX, 19)				-0.177 (-1.13)	0.993** (2.33)
Imp. Fin. (DEX, 8)				-0.249 (-1.19)	0.011 (0.05)
Inf. Launch (DEX, 3)				-1.003*** (-2.97)	-0.654** (-2.06)
Infinity Pad (DEX, 9)				-0.902*** (-8.19)	-0.660*** (-5.49)
KickPAD (DEX, 7)				-0.821*** (-5.35)	-0.480*** (-3.10)
Kommunitas (DEX, 9)				-1.065*** (-5.71)	-0.814*** (-4.11)
Krystal GO (DEX, 13)				-0.707*** (-4.50)	-0.474*** (-2.92)
Kucoin (CEX, 13)				-0.426** (-2.23)	0.789* (1.79)
LaunchZone (DEX, 16)				-0.468*** (-2.86)	-0.225 (-1.32)
Launchpool (DEX, 12)				-0.235 (-1.07)	0.046 (0.20)
Lemonade (DEX, 2)				-0.471*** (-4.30)	-0.136 (-0.96)
Lightning (DEX, 8)				0.166 (1.27)	0.496*** (4.24)
MEXC (CEX, 10)				-0.972*** (-5.18)	0.207 (0.47)
MISO (DEX, 8)				1.757*** (3.83)	2.259*** (4.61)
MoonEdge (DEX, 4)				-0.712*** (-3.92)	-0.410** (-2.40)
MoonStarter (DEX, 19)				-0.955*** (-7.11)	-0.680*** (-5.47)
OKX (CEX, 1)				0.778*** (6.56)	2.038*** (4.76)
OccamRazer (DEX, 17)				-0.416*** (-2.95)	-0.143 (-0.97)
Other (203)				0.000	0.000
PAID (DEX, 64)				-0.434*** (-3.70)	-0.155 (-1.34)
PancakeSwap (DEX, 14)				0.888*** (4.50)	1.374*** (7.23)
Polkabridge (DEX, 11)				-1.371*** (-6.24)	-1.109*** (-4.96)
Polkastarter (DEX, 88)				-0.425*** (-3.65)	-0.073 (-0.62)
Poolz (DEX, 61)				-0.797*** (-7.21)	-0.496*** (-4.69)
Probit (CEX, 4)				-0.791* (-1.88)	0.874* (1.71)
Red Kite (DEX, 50)				-0.780*** (-7.44)	-0.486*** (-4.61)
Republic (DEX, 5)				2.693*** (8.90)	2.265*** (6.14)
Roseon Fin. (DEX, 11)				-1.013*** (-6.03)	-0.762*** (-4.72)
Scaleswap (DEX, 7)				-0.917*** (-4.26)	-0.667*** (-2.68)
Seedify (DEX, 34)				-0.185* (-1.81)	0.090 (0.88)
SolRazr (DEX, 8)				0.416 (1.45)	0.668** (2.19)
Solanium (DEX, 22)				-0.554*** (-3.27)	-0.294* (-1.71)
StarLaunch (DEX, 3)				0.015 (0.07)	0.276 (1.48)
Starter (DEX, 18)				-0.890*** (-5.50)	-0.647*** (-3.94)
Synapse (DEX, 15)				-0.730*** (-3.95)	-0.518*** (-2.80)
TronPad (DEX, 7)				-0.741*** (-7.33)	-0.511*** (-5.06)

	(i) Basic	(ii) Type	(iii) Category	(iv) Platform	(v) Full
TruePNL (DEX, 21)				-1.345*** (-8.92)	-1.143*** (-7.10)
TrustPad (DEX, 55)				-0.712*** (-6.55)	-0.464*** (-4.22)
TrustSwap (DEX, 7)				-0.491* (-1.71)	-0.314 (-1.05)
Vent Fin. (DEX, 1)				-0.525*** (-5.51)	-0.269*** (-2.84)
WeStarter (DEX, 11)				-2.120*** (-11.31)	-1.802*** (-9.40)
Yellow Rd. (DEX, 5)				-1.398*** (-5.16)	-0.979*** (-4.48)
ZB (CEX, 1)				-1.412*** (-15.19)	-0.228 (-0.55)
ZENDIT (DEX, 34)				-0.902*** (-6.06)	-0.580*** (-3.91)
ZeeDO (DEX, 14)				-1.034*** (-6.68)	-0.726*** (-4.86)
Zel. Fin. (DEX, 1)				-1.312*** (-11.23)	-1.356*** (-6.93)

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Parameters for the linear regression model with logRaised as the dependent variable. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 ICOs. The model specifications are: (i) Basic, followed by additional specifications controlling for the token offering's type (ii), category (iii) and token offering platform (iv), and also a full specification (v) which contains all controls. Note that control variable names include in brackets the number of offerings in each case and, for the platform controls, whether the launchpad platform is a centralized (CEX) or decentralized (DEX) exchange. The first column in each model specification displays the regression betas and the second column displays the t-statistics in parentheses.

Regarding the explanatory variables included in the Basic configuration shown in Table 4.10, we observe a significant positive effect of logtCap on fundraising success, consistent with our hypothesis ($H1$) that the target market cap can predict the actual amount raised in the offering. Interestingly, we further observe an increase in the magnitude of the regression coefficient for logtCap to 0.44, compared with the corresponding results from the 2017 – 2020 period with an average effect of 0.33, which suggests that logtCap is now an even more significant predictor of fundraising success in comparison with previous stages in the token offering market. Similarly, the ETH price variable is found to have a positive effect, consistent with the earlier 2017 – 2019 periods and in contrast with the negative effect observed during the 2020 pandemic period; this is in line with our hypothesis $H3$ that the average ETH/USD price level during a token offering has a generally positive effect on fundraising success. Surprisingly and contrary to our hypothesis ($H3$), the average 30-day momentum of ETH/USD returns during each offering have a significant negative effect on logRaised, which could imply overall that the optimal timing for launching a token offering

in 2021 - 2022 was during the ‘cooling-off’ periods of negative momentum in the ETH/USD returns. Additionally, we observe that the combined volume of Google web searches, Google news searches and YouTube searches on the ‘Bitcoin’, ‘Ethereum’ and ‘Blockchain’ keywords has no significant effect on logRaised, indicating that investors are no longer influenced by the overall investor attention to the crypto asset space.


As expected, the number of token offerings launched by the same venture has a negative effect on the amount raised, as the total funding required for the venture is achieved from multiple offerings. Finally, the effect of distributing the token’s supply shares more equally is relatively small and only significant in the Basic model configuration, and the duration of the token offering now exhibits a significant positive effect, contrary to our findings in the 2017 - 2020 period.

To further illustrate the additional effects of the Type, Category and Platform controls included in models (ii) - (iv) of Table 4.10, we further examine their coefficients added to the overall constant coefficient in each specification. This provides an intuitive measure of comparison between the different types, categories and platforms, in the form of the net offset in standard deviation units of the dependent variable. In the Type model specification (ii), the constant coefficient shown in Table 4.10 is 1.979; we choose the ICO type as base so its coefficient is zero and the IDO and IEO coefficients are -1.731 and -2.2 respectively. The resulting net offsets are therefore 1.979 for ICOs, 0.248 for IDOs and -0.221 for IEOs, while the sample includes 54 ICOs, 1,340 IDOs and 213 IEOs. This suggests that although the choice of conducting an ICO is no longer very common, these offerings tend to aim for and raise more funding; similarly, IDOs are more prevalent in the sample and although their offset is smaller than that of ICOs, it is still positive, indicating that it is still possible to raise significant amounts of funding via IDOs; finally, for IEOs the effect on logRaised is significantly negative, so that the net offset is also negative, suggesting that IEOs are perhaps considered by token issuers more as a means of

on-boarding their token to a crypto asset exchange and perhaps as a marketing tool, rather than the primary means of fundraising.

We perform a similar analysis for the token offering platforms, showing their net regression coefficient offsets in Figure C1 and the number of token offerings per platform in Figure C2, both included in Appendix C, and show the combined scatter plot in Figure 4.2. We see that the largest net offsets are achieved by relatively less popular platforms such as Republic, MISO and Gnosis, while Gate.io which is the most popular platform has one of the most negative net offsets. Given that a larger number of token offerings in an IDO platform could attract more investors and provide more liquidity at the trading stage, it may be advisable for token issuers to compromise and list their offering in platforms such as PancakeSwap or Avalaunch, which have a slightly lower net offset in terms of the funding raised but are somewhat more popular.

Figure 4.2: Token offerings per launchpad platform and coefficients




Scatter plot of the constant coefficients for the token offering launchpad platform based on the linear regression model with $\log(\text{Raised})$ as the dependent variable, against the number of token offerings per platform shown in a logarithmic axis. Different colours are used to denote whether the launchpad platform is a centralized (CEX in red) or decentralized (DEX in blue) exchange. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.

Finally, we rank-order the token offering categories by their corresponding net regression coefficient offsets as shown in Figure C3 and show the number of token offerings per category in Figure C4 – both included in Appendix C, and combine

these two metrics in a scatter plot in Figure 4.3. The largest net offsets are observed for categories such as Energy, Artificial (AI) and Data which only include 25 token offerings in total, suggesting that less competition in the token offering category allows the token issuers to capture more of the available funding in that space. At the same time, very popular categories such as Gaming, DeFi, Blockchain and Non-Fungible (NFT-related ventures) still have a positive net offset, indicating that considerable funding is available in these spaces as well.

Figure 4.3: Token offerings per category and coefficients



Scatter plot of the constant coefficients for the token offering category based on the linear regression model with logRaised as the dependent variable, against the number of token offerings per category shown in a logarithmic axis. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.

4.6 Summary and Conclusions

In this Chapter we present empirical evidence on factors that influence the fundraising success of token offerings between 2017 – early 2022. Following the relevant literature, we proxy success with the amount of funding raised and – where possible, the exceedance of the soft cap. Possible success determinants are selected based on the offering characteristics, the token’s structure, the token issuers’ public profile and online presence, and also indicative variables for the entire token offerings space and crypto asset market.

The chapter's incremental contributions are as follows. We examine the entire 2017 – early 2022 period, which – at the time of writing, constitutes the most comprehensive study on the subject in terms of sample period coverage. We introduce success determinants that were not used previously in the relevant literature, such as the token's target market capitalization and the launchpad platform used by offerings in 2021 – early 2022. We discuss different perspectives on variables such as the token offering rating and tax-haven domicile; and we revisit variables for which there are conflicting findings in the relevant literature – such as the token presale.

We obtain relevant data from ICObench and Cryptorank and use other ranking and rating websites as supplementary data sources; the final sample size comprises 2,926 offerings, again constituting one the most comprehensive studies in this respect. We compare and extend the findings in the relevant literature for the 2017 – 2019 period, with a separate examination of the earlier 2017 – mid-2018 period of significant activity, and the subsequent mid-2018 – 2019 period of declining interest. We identify 2020 as a period of significantly reduced activity, due to the initial peak of the Covid-19 pandemic.

Crucially, this Chapter is the first relevant research to identify a revival in significant investor interest for token offerings during 2021 – early 2022, with new all-time highs in the number of token offerings completed per month in late 2021 and early 2022 and a significant shift with the emergence and near-complete domination of initial decentralized exchange offerings (IDOs). Based on these market changes, our findings for the 2021 – early 2022 period demonstrate a shift in the fundraising success determinants, for instance with the launchpad platform used playing a significant role in explaining token offering fundraising success.

We form hypotheses on the expected effect of several factors on fundraising success: a higher target market cap has an expected positive effect on fundraising success; a higher soft cap has a negative effect on the probability of soft cap exceedance; signals of a venture's public credibility have a positive effect on fundraising success but only if they are genuine; increased investor attention and belief in the crypto asset class,

as well as increased hype have a positive effect; a tax-haven or undisclosed venture domicile has a positive effect; bonus schemes, token distribution and the offering's duration have no impact on fundraising success.

Having discussed the chapter's contributions and research hypotheses, we end this Section with a summary of the findings for the 2017 – 2020 and 2021 – early 2022 sample periods examined, grouped by hypothesis.

Consistent with our hypothesis (*H1*) that a project's target market cap is a good predictor of the actual amount raised in the token sale, the logtCap coefficient is significant and positive in all sample periods and model configurations examined; additionally, we find a significant negative coefficient for logSoftcap in the probit model for the 2017 – 2020 period, again consistent with our suggestion in *H1*.

Our findings are mixed and not always consistent with our assumption (*H2*) that fundraising success only increases with genuine signals of public credibility; more specifically, we identify Rating and Presale as genuine signals and argue that Team and Advisors can be considered as cheap signals. Interestingly, Rating has a significant positive effect in the 2017 – mid-2018 period but *not* in the subsequent mid-2018 - 2020 periods, indicating that investors may be influenced by the allegations of rating shopping, and the warnings of moral hazard and potentially skewed incentive structures for rating websites, as highlighted in the relevant literature. On the contrary, the effect of holding a token presale prior to the public offering has no effect on fundraising success; a higher team size has a mostly positive effect, while the advisory team size is insignificant in all cases. It therefore appears that token offering investors may not be consistent in distinguishing between genuine and cheap credibility signals, which confirms previous findings in the relevant literature.

Additionally, in the 2021 – early 2022 sample period, the model specification that includes the launchpad platform control variable exhibits a very high R^2 value, suggesting that token offering launchpad platforms are now key intermediaries in the ecosystem, and prospective investors possibly consider token offerings conducted on established and reliable launchpad platforms as more credible.

Regarding $H3$, the ETH price variable – i.e. the average ETH/USD price during each offering, has a mostly positive effect on success during the 2017 – 2019 period, indicating that it does capture investor trust in the entire crypto asset class. However, the ETH price coefficient is negative in the linear model during the 2020 pandemic peak period and it is insignificant in the probit model. During the 2021 – early 2022 period, the ETH price variable is again found to have a positive effect, consistent with the assumption of $H3$. However, the ETH mom variable – i.e. the ETH/USD average momentum, is insignificant in most cases and even has a negative effect in the 2021 - early 2022 period. Additionally, it appears that in the 2021 – early 2022 period, investors are no longer influenced by the overall investor attention to the crypto asset space, as proxied by the combined volume of Google web searches, Google news searches and YouTube searches for the ‘Bitcoin’, ‘Ethereum’ and ‘Blockchain’ terms.

Our hypothesis ($H4$) on investor preference for tax-haven-domiciled ventures is partially confirmed, in that the OFC (offshore financial centre) variable has a mostly positive effect in the 2017 – 2019 period, but not in the 2020 pandemic period. Moreover, regarding the most recent 2021 – early 2022 period, it is possible that venture domicile is less important to prospective investors, compared e.g. to the launchpad platform’s domicile.

Our findings on the coefficients of the Bonus, Distribution and logDuration variables show that they are important in the early 2017 – mid-2018 period, less important in the mid-2018 – 2019 period and insignificant during the 2020 pandemic period; this partially confirms our hypothesis ($H5$) that they have no effect on fundraising success. Interestingly, the sign of the Distribution variable’s coefficient flips from negative in the 2017 – mid-2018, to positive in mid-2018 – 2019, so that the suggestion of Ahlers et al. (2015) on equity crowdfunding that higher equity retention by the venture should reduce investor uncertainty and increase fundraising success is only partially confirmed for token offerings.

Finally, note that due to the significant differences in the characteristics of token offerings in the 2021 – early 2022 period, the corresponding linear model configuration

explaining fundraising success uses a somewhat different set of explanatory variables. In fact, the 2021 – early 2022 sample results show that the model specification including the launchpad platform control variable exhibits a very high R^2 value, compared against both earlier sample periods and results in the relevant literature. This suggests that the launchpad platform used to conduct a token offering – together with the type of offering, prove to be the most significant determinants of token offering fundraising success in 2021 – early 2022.

SUMMARY AND CONCLUSIONS

The ever-changing landscape of crypto asset markets poses several significant challenges. For instance, the extensive use of crypto assets as a speculative instrument and the significantly faster pace compared with trading in traditional financial assets creates an environment of highly volatile prices, so that market risk measurement is both highly important and technically challenging. The lack of a regulatory framework coupled with the significant operational risks in crypto asset investments render market manipulation relatively easy to perform and even easier to get away with. Also, the area of crowdfunding via token offerings has undergone significant changes since 2017 when this space gained investor attention for the first time, so the characteristics that render such ventures successful are quite different in 2022 compared with 2017. Challenges such as the above have attracted significant academic interest and formed a cutting-edge and very active research area, with more than 2,200 relevant articles published since 2013 in finance, economics, econometrics, business, management and accounting journals.

This doctoral thesis develops three distinct empirical finance topics, which study and attempt to address the challenges highlighted above. Chapter 2 focuses on the modelling of crypto asset market risk; Chapter 3 examines the potential manipulation of crypto asset markets via stablecoins – tether in particular; and Chapter 4 studies the fundraising success factors of crowdfunding via token offerings. In the following, we provide a brief summary of the relevant findings and conclusions, ending with an outlook on potential future developments in academic research for crypto asset and blockchain finance.

Crypto Asset Market Risk Modelling

In Chapter 2, we conduct an out-of-sample volatility, Value-at-Risk and Expected Shortfall analysis on the daily and hourly log returns of the large-cap crypto assets bitcoin, ether, ripple and litecoin. We model volatility with a range of simple, less simple and complex models and examine particularly how simpler models with ad hoc parameter choices perform against more complex models with regard to their forecasting accuracy. The simplest model used is an equally-weighted moving average model, which we use as a benchmark; we also use RiskMetrics-type EWMA models, further introducing an asymmetric EWMA specification similar to an AGARCH model but with ad hoc parameter choices; the forecasting performance of these simpler models is compared against the more complex symmetric and asymmetric GARCH models. Beyond volatility and quantile risk measure forecasts, the above models are extended in a multivariate setting via the multivariate random walk, EWMA and AEWMA models, and the symmetric and asymmetric DCC models to produce forecasts of the entire covariance matrix of bitcoin, ether, ripple and litecoin returns.

Daily forecasts are produced between January 2017 – August 2021 and hourly forecasts between 1 May 2021 – 1 July 2021. The 1-period-ahead left- and right-tail Value-at-Risk and Expected Shortfall are forecasted for each asset and at multiple significance levels. VaR forecasts are backtested using the conditional coverage and the industry standard traffic light tests and similarly, ES is backtested with the exceedance residual test and again a modified traffic light test, for which we develop an extension to right-tail ES forecasts. The accuracy of volatility and covariance forecasts is also assessed using the continuous ranked probability score (CRPS) for each asset and multivariate scoring rules such as the energy and variogram scores for the joint density forecasts of bitcoin, ether, ripple and litecoin returns.

The results for daily and hourly left- and right-rail VaR and ES forecasts show that the asymmetric EWMA models with ad hoc parameter choices are almost always just as accurate as a more complex EGARCH specification. In fact, there are

cases when the simpler RiskMetrics-type symmetric EWMA specifications and even the equally-weighted moving average random walk model are found to be on par with the EGARCH model. As for volatility forecasts, the test of equal forecasting performance based on CRPS and multivariate scores shows that all models examined, even the random walk benchmark, are equally accurate, both at the daily and hourly frequencies. This important result extends previous findings in the relevant literature that simpler models can sometimes outperform more complex ones in terms of volatility and quantile risk measure forecasting accuracy, so the process of crypto asset market risk measurement could be simplified significantly for practitioners.

Blockchain Transactions and Crypto Asset Market Manipulation

In Chapter 3 we examine the potential manipulation of crypto asset markets – particularly bitcoin, via large blockchain transactions of the tether stablecoin, originating from the centralized tether treasury. The methodology introduced is adapted from the framework of securities fraud litigation event studies, and determines the materiality of price effects based on regression factor models, estimating abnormal bitcoin returns during each USDT outflow event from the tether treasury.

The contributions of this chapter are as follows: an event study-based methodology is developed for the determination of materiality and estimation of damages incurred by potential market manipulation, which is a necessary component in relevant legal cases, as damage estimates in securities fraud cases are generally not accepted in court unless an event study-based approach is employed. The methodology developed is considered an improvement compared with similar research examined in the literature review, for instance by adopting a flexible-length event window approach to account for the significant clustering of tether outflow events, also controlling for the effect of past events. Additionally, based on the experience of deploying a blockchain node locally, significant detail is provided on the retrieval process of blockchain transaction data, for which the retrieval techniques are sparsely covered in the relevant literature.

The main results are based on an equally-weighted market portfolio index model specification, while results from alternative model specifications are also provided. The overall findings indicate that between April 2019 – January 2021, 30 out of 671 USDT outflow events from the tether treasury towards crypto asset exchanges and other unknown entities produce a significant positive abnormal bitcoin return and therefore a material positive price impact indicative of price manipulation. The total damages from the above material events are estimated at \$54 million, when by comparison, a legal case involving client and market deception resulted in Bitfinex and Tether paying \$18.5 million in fines. The indication of particular tether blockchain transactions as having a significant effect on bitcoin prices is partially consistent with the findings in the relevant literature, particularly for articles which employ an event study-based methodology.

The Tokenomics of Crowdfunding

Chapter 4 presents evidence on the fundraising success factors of token offerings between 2017 – early 2022. Fundraising success is proxied with the amount of funding raised and – where possible, the exceedance of the minimum funding target. Success factors are examined using linear and probit regression models and are selected based on the token offering characteristics, the venture’s online presence, and also representative variables for the entire crypto asset and token offerings market. We have examined the entire 2017 – early 2022 period with a sample of 2,926 offerings, which – at the time of writing, constitutes one of the most comprehensive studies in the area of token offering fundraising success.

We identify a revival in token offering activity during the 2021 – early 2022 period and also the emergence of initial decentralized exchange offerings (IDOs). The chapter’s hypotheses relate to the expected effect of potential factors on fundraising success: a positive effect for: a higher target market cap, a lower soft cap, genuine signals of venture public credibility, increased hype and investor attention and belief in crypto assets, and a tax-haven venture domicile. Additionally, the existence

of bonus schemes, the token’s supply distribution and the offering’s duration are hypothesized to have no impact on fundraising success.

Our findings indicate that a token’s target market has significant explanatory power, consistent with our hypothesis. Regarding signals of venture public credibility, it appears that investors do not distinguish consistently between genuine and cheap signals, which also confirms previous findings in the relevant literature. We also find that token offering launchpad platforms in 2021 – early 2022 have become key intermediaries, and investors perhaps interpret the use of established and reliable launchpad platforms as a credibility signal for token offerings. Regarding investor attention and belief, we find that the price of ether possibly captures investor trust in the entire crypto asset class, but the results are not significant for ether’s momentum and Google search volume.

Tax-haven-domiciled ventures are found to be more successful in fundraising during the 2017 – 2019 period but not in 2020; additionally, as the majority of token offerings in the 2021 – early 2022 period are issued on launchpad platforms, it is possible that venture domicile is less important to prospective investors, compared with the launchpad platform’s domicile. Finally, our findings on the effect of a bonus scheme, the token’s supply distribution and the offering’s duration show that they are somewhat significant in 2017 – 2019, but insignificant in 2020. Specifically for the distribution variable, this indicates that the argument of investor uncertainty reduction via higher equity retention does not hold consistently for token offerings.

Outlook

The role of the blockchain-enabled Web 3.0 has introduced ‘a parallel system of finance of unprecedented flexibility and creativity in less than a decade’ (Weyl et al., 2022). Even in the relatively short lifetime of crypto assets, multiple innovative concepts have already been introduced: after the first wave of payment-type cryptocurrencies such as bitcoin and litecoin, there came smart contract platforms such as Ethereum, the surge in crowdfunding via token offerings, the development of several stablecoins,

the decentralized finance (DeFi) wave with decentralized exchanges (DEXs) and crypto asset lending platforms, and also the increased interest in digital collectibles markets and metaverse-related applications made possible via non-fungible tokens (NFTs) and relevant marketplaces. Crypto asset markets are in a continuous state of flux, with ongoing developments both on the innovation side and on the regulatory side. Future research in crypto asset and blockchain finance could therefore focus – and is already beginning to focus, on such developments.

Similarly, regarding the topics developed in this doctoral thesis, further research could again be motivated by ongoing and future developments, and could also focus on features of these topics that were not developed as part of this thesis. For instance, in the area of crypto asset market risk, crypto asset volatility could in the future exhibit very different characteristics, such as becoming much lower after stricter regulatory frameworks are in place; this potential development is not unrealistic given upcoming regulatory framework updates, such as the European Union-related Transfer of Funds Regulation (TFR) and Markets in Crypto Assets (MiCA) rules still under discussion by the European Council, which may take effect over the coming 18 months (Chainalysis, 2022). In the area of market manipulation and stablecoins, a study of the recent Terra USD/Luna meltdown which caused a downward spiral in crypto asset markets (Shen, 2022) could potentially identify drivers of systematic risk in crypto asset markets. Finally, in the area of crowdfunding via token offerings, several studies examine the market performance of such tokens once these are listed, but as this space is rife with scams and various types of rug pulls (Puggioni, 2022), models that analyse such cases and could potentially estimate the probability of such rug pulls would be of significant academic interest and also of significant value to crypto asset investors.

Finally, ongoing research has the potential to generalize the concept of a blockchain even further. For instance, Weyl et al. (2022) introduce the concept of ‘soulbound’ tokens which enable the creation of a ‘decentralized society’ (DeSoc), a pluralistic ecosystem in which communities come together bottom-up, as emergent properties

of each other to co-create plural network goods and intelligences. At the risk of sounding cliché, the sky is the limit.

BIBLIOGRAPHY

- Aas, K. and Haff, I. H. (2006), ‘The generalized hyperbolic skew Student-t distribution’, *Journal of Financial Econometrics* **4**(2), 275–309.
- Acerbi, C. and Tasche, D. (2002), ‘Expected shortfall: A natural coherent alternative to value at risk’, *Economic Notes* **31**(2), 379–388.
- Acereda, B., Leon, A. and Mora, J. (2020), ‘Estimating the expected shortfall of cryptocurrencies: An evaluation based on backtesting’, *Finance Research Letters* **33**, 101181.
- Adhami, S., Giudici, G. and Martinazzi, S. (2018), ‘Why do businesses go crypto? An empirical analysis of initial coin offerings’, *Journal of Economics and Business* **100**, 64–75.
- Aggarwal, R., Hanley, K. W. and Zhao, X. (2019), ‘The role of information production in private markets: Evidence from initial coin offerings’. Semantic Scholar working paper 204959864.
- Ahlers, G. K., Cumming, D., Günther, C. and Schweizer, D. (2015), ‘Signaling in equity crowdfunding’, *Entrepreneurship Theory and Practice* **39**(4), 955–980.
- Aielli, G. P. (2013), ‘Dynamic Conditional Correlation: On properties and estimation’, *Journal of Business & Economic Statistics* **31**(3), 282–299.
- Akerlof, G. (1970), ‘The market for ‘lemons’: Quality uncertainty and the market mechanism’, *The Quarterly Journal of Economics* **84**(3), 488–500.
- Al-Khazali, O., Bouri, E. and Roubaud, D. (2018), ‘The impact of positive and negative macroeconomic news surprises: Gold versus bitcoin’, *Economics Bulletin* **38**(1), 373–382.
- Al-Yahyaee, K., Mensi, W., Al-Jarrah, I., Hamdi, A. and Kang, S. (2019), ‘Volatility forecasting, downside risk, and diversification benefits of bitcoin and oil and international commodity markets: A comparative analysis with yellow metal’, *North American Journal of Economics and Finance* **49**, 104–120.
- Albrecht, S., Lutz, B. and Neumann, D. (2020), ‘The behavior of blockchain ventures on Twitter as a determinant for funding success’, *Electronic Markets* **30**(2), 241–257.
- Alexander, C. and Dakos, M. (2019), ‘Bitcoin: welcome to bitfinex’s second tether bubble’, Medium.
- Alexander, C. and Dakos, M. (2020), ‘A critical investigation of cryptocurrency data and analysis’, *Quantitative Finance* **20**(2), 173–188.
- Alexander, C. and Heck, D. F. (2020), ‘Price discovery in bitcoin: The impact of unregulated markets’, *Journal of Financial Stability* **50**, 100776.

- Alstadsæter, A., Johannesen, N. and Zucman, G. (2018), ‘Who owns the wealth in tax havens? Macro evidence and implications for global inequality’, *Journal of Public Economics* **162**, 89–100.
- Amsden, R. and Schweizer, D. (2018), ‘Are blockchain crowdsales the new ‘gold rush’? Success determinants of initial coin offerings’. SSRN working paper 3163849.
- Ante, L. and Fiedler, I. (2020), ‘Cheap signals in security token offerings (STOs)’, *Quantitative Finance and Economics* **4**(QFE-04-04-028), 608.
- Ante, L., Fiedler, I. and Strehle, E. (2021a), ‘The impact of transparent money flows: Effects of stablecoin transfers on the returns and trading volume of Bitcoin’, *Technological Forecasting and Social Change* **170**.
- Ante, L., Fiedler, I. and Strehle, E. (2021b), ‘The influence of stablecoin issuances on cryptocurrency markets’, *Finance Research Letters* **41**.
- Antonopoulos, A. M. (2016), *The internet of money: Talks by Andreas M. Antonopoulos*, Merkle Bloom LLC.
- Antonopoulos, A. M. (2017), *Mastering Bitcoin: Programming the open blockchain*, O’Reilly Media, Inc.
- Ardia, D., Bluteau, K. and Rüede, M. (2019), ‘Regime changes in bitcoin GARCH volatility dynamics’, *Finance Research Letters* **29**, 266–271.
- Ardia, D., Boudt, K. and Gagnon-Fleury, J.-P. (2017), ‘RiskPortfolios: Computation of risk-based portfolios in R’, *Journal of Open Source Software* **2**(10), 171.
- Armitage, S. (1995), ‘Event study methods and evidence on their performance’, *Journal of Economic Surveys* **9**(1), 25–52.
- Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D. (1999), ‘Coherent measures of risk’, *Mathematical Finance* **9**(3), 203–228.
- Baker, A. C. (2016), ‘Single-firm event studies, securities fraud, and financial crisis: problems of inference’, *Stan. L. Rev.* **68**, 1207–1266.
- Baker McKenzie (2018), ‘Initial coin offering: A legal work plan’. Report available at: <https://bit.ly/36lhNtM>.
- Barucci, E., Marazzina, D. and Giuffra Moncayo, G. (2021), ‘Cryptocurrencies and stablecoins: a high frequency analysis’. SSRN working paper 4018171.
- Basel Committee (1996), ‘Supervisory framework for the use of backtesting in conjunction with the internal models approach to market risk capital requirements’, Basel Committee on Banking Supervision. Available at: <https://www.bis.org/publ/bcbs22.htm>.
- Basel Committee (2012), ‘Fundamental review of the trading book: A revised market risk framework’, Basel Committee on Banking Supervision. Available at: <https://www.bis.org/publ/bcbs219.htm>.
- Baum, J. A. and Silverman, B. S. (2004), ‘Picking winners or building them? Alliance, intellectual, and human capital as selection criteria in venture financing and performance of biotechnology startups’, *Journal of Business Venturing* **19**(3), 411 – 436.

- Baur, D., Dimpfl, T. and Kuck, K. (2018), ‘Bitcoin, gold and the US dollar: A replication and extension’, *Finance Research Letters* **25**, 103–110.
- Bauwens, L. and Storti, G. (2009), ‘A component GARCH model with time varying weights’, *Studies in Nonlinear Dynamics & Econometrics* **13**(2), 1–33.
- Bayer, S. and Dimitriadis, T. (2020), ‘Regression-based expected shortfall backtesting’, *Journal of Financial Econometrics* . nbaa013.
- Bazán-Palomino, W. (2020), *Bitcoin and its offspring: A volatility risk approach*, Springer Singapore, pp. 233–256.
- Belitski, M. and Boreiko, D. (2021), ‘Success factors of initial coin offerings’, *Journal of Technology Transfer* . Article in Press.
- Blaseg, D. (2018), ‘Dynamics of voluntary disclosure in the unregulated market for initial coin offerings’. SSRN working paper 3207641.
- Bollerslev, T. (1986), ‘Generalized autoregressive conditional heteroskedasticity’, *Journal of Econometrics* **31**(3), 307–327.
- Bonello, A. and Suda, D. (2018), Volatility regime analysis of bitcoin price dynamics using Markov-switching GARCH models, *in* ‘Proceedings of the European Simulation and Modelling Conference 2018’, pp. 213–218.
- Boreiko, D. and Risteski, D. (2021), ‘Serial and large investors in initial coin offerings’, *Small Business Economics* **57**(2), 1053–1071.
- Boreiko, D. and Vidusso, G. (2019), ‘New blockchain intermediaries: Do ICO rating websites do their job well?’, *Journal of Alternative Investments* **21**(4), 67–79.
- Bouoiyour, J. and Selmi, R. (2014), ‘Commodity price uncertainty and manufactured exports in Morocco and Tunisia: Some insights from a novel GARCH model’, *Economics Bulletin* **34**(1), 220–233.
- Bouoiyour, J. and Selmi, R. (2016), ‘Bitcoin: A beginning of a new phase?’, *Economics Bulletin* **36**(3), 1430–1440.
- Bouri, E., Molnar, P., Azzi, G., Roubaud, D. and Hagfors, L. (2017), ‘On the hedge and safe haven properties of bitcoin: Is it really more than a diversifier?’, *Finance Research Letters* **20**, 192–198.
- Bourveau, T., De George, E., Ellahie, A. and Macciocchi, D. (2022), ‘The role of disclosure and information intermediaries in an unregulated capital market: Evidence from initial coin offerings’, *Journal of Accounting Research* **60**(1), 129–167.
- Bovaird, C. (2021), ‘Bitcoin price volatility reached a 14-month high in June’, Forbes. Available at: <https://bit.ly/3zLZtXw>.
- Bovet, A., Campajola, C., Lazo, J. F., Mottes, F., Pozzana, I., Restocchi, V., Saggese, P., Vallarano, N., Squartini, T. and Tessone, C. J. (2018), ‘Network-based indicators of bitcoin bubbles’, *arXiv preprint arXiv:1805.04460* .
- Boydston, A. E., Bevan, S. and Thomas III, H. F. (2014), ‘The importance of attention diversity and how to measure it’, *Policy Studies Journal* **42**(2), 173–196.

- Brown, S. J. and Warner, J. B. (1985), ‘Using daily stock returns: The case of event studies’, *Journal of Financial Economics* **14**(1), 3–31.
- Burniske, C. and Tatar, J. (2018), *Cryptoassets: The innovative investor’s guide to bitcoin and beyond*, McGraw-Hill Education.
- Burns, L. and Moro, A. (2018), ‘What makes an ICO successful? An investigation of the role of ICO characteristics, team quality and market sentiment’. SSRN working paper 3256512.
- Cameron, A. C. and Trivedi, P. K. (2009), *Microeconometrics using Stata*, Vol. 5, College Station, TX: Stata press.
- Campino, J., Brochado, A. and Rosa, A. (2022), ‘Initial coin offerings (ICOs): Why do they succeed?’, *Financial Innovation* **8**(1), 1–35.
- Caporale, G. and Zekokh, T. (2019), ‘Modelling volatility of cryptocurrencies using Markov-switching GARCH models’, *Research in International Business and Finance* **48**, 143–155.
- Cappiello, L., Engle, R. F. and Sheppard, K. (2006), ‘Asymmetric dynamics in the correlations of global equity and bond returns’, *Journal of Financial Econometrics* **4**(4), 537–572.
- Carter, N. (2018), ‘The dark underbelly of cryptocurrency markets’, Medium.
- Catania, L. and Grassi, S. (2021), ‘Forecasting cryptocurrency volatility’, *International Journal of Forecasting* . Article in press.
- Catania, L., Grassi, S. and Ravazzolo, F. (2019), ‘Forecasting cryptocurrencies under model and parameter instability’, *International Journal of Forecasting* **35**(2), 485–501.
- Chae, J. (2005), ‘Trading volume, information asymmetry, and timing information’, *The Journal of Finance* **60**(1), 413–442.
- Chainalysis (2022), ‘EU agrees on comprehensive cryptocurrency regulations: Part 1 of our breakdown’, Chainalysis.
- Chen, K. (2019), ‘Information asymmetry in initial coin offerings (ICOs): Investigating the effects of multiple channel signals’, *Electronic Commerce Research and Applications* **36**, 100858.
- Christoffersen, P. (1998), ‘Evaluating interval forecasts’, *International Economic Review* **39**(4), 841–862.
- Chu, J., Chan, S., Nadarajah, S. and Osterrieder, J. (2017), ‘GARCH modelling of cryptocurrencies’, *Journal of Risk and Financial Management* **10**(4), 17.
- Clift, S., Costanzino, N. and Curran, M. (2016), ‘Empirical performance of backtesting methods for expected shortfall’, SSRN Discussion Paper 2618345.
- Connelly, B. L., Certo, S. T., Ireland, R. D. and Reutzel, C. R. (2011), ‘Signaling theory: A review and assessment’, *Journal of Management* **37**(1), 39–67.
- Corbet, S., Lucey, B., Peat, M. and Vigne, S. (2018), ‘Bitcoin futures: What use are they?’, *Economics Letters* **172**, 23–27.

- Corrado, C. J. (2011), ‘Event studies: A methodology review’, *Accounting & Finance* **51**(1), 207–234.
- Costanzino, N. and Curran, M. (2015), ‘Backtesting general spectral risk measures with application to expected shortfall’, SSRN Discussion Paper 2514403.
- Costanzino, N. and Curran, M. (2018), ‘A simple traffic light approach to backtesting expected shortfall’, *Risks* **6**(1), 2.
- Creal, D., Koopman, S. J. and Lucas, A. (2013), ‘Generalized autoregressive score models with applications’, *Journal of Applied Econometrics* **28**(5), 777–795.
- Cryptocompare (2018), ‘Cryptoasset taxonomy report’, Cryptocompare Research.
- Czaja, D. and Röder, F. (2021), ‘Signalling in initial coin offerings: The key role of entrepreneurs’ self-efficacy and media presence’, *Abacus* . Article in Press.
- Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L. and Juels, A. (2019), ‘Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges’, arXiv.
- Dakos, M. (2020), ‘Bitcoin whales and crypto market manipulation’, Medium.
- Daye, W., Guan, C., Swift, L. and Cheema, S. (2019), ‘Understanding the crypto asset market’, Quantifi. Available at: quantifisolutions.com/understanding-the-cryptoasset-market.
- Diebold, F. X. and Mariano, R. S. (1995), ‘Comparing predictive accuracy’, *Journal of Business and Economic Statistics* **13**(3), 253–263.
- Ding, Z., Granger, C. and Engle, R. (1993), ‘A long memory property of stock market returns and a new model’, *Journal of Empirical Finance* **1**(1), 83–106.
- Dove, T., Heath, D. and Heaton, J. (2019), ‘Bias-corrected estimation of price impact in securities litigation’, *American Law and Economics Review* **21**(1), 184–208.
- Dyhrberg, A. (2016), ‘Bitcoin, gold and the dollar - A GARCH volatility analysis’, *Finance Research Letters* **16**, 85–92.
- Easley, D., O’Hara, M. and Basu, S. (2019), ‘From mining to markets: The evolution of bitcoin transaction fees’, *Journal of Financial Economics* **134**(1), 91–109.
- Engle, R. (2002), ‘Dynamic Conditional Correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models’, *Journal of Business & Economic Statistics* **20**(3), 339–350.
- Engle, R. F., Focardi, S. M. and Fabozzi, F. J. (2008), *ARCH/GARCH models in applied financial econometrics*, John Wiley and Sons, chapter 60.
- Engle, R. F. and Kroner, K. F. (1995), ‘Multivariate simultaneous generalized ARCH’, *Econometric Theory* **11**(1), 122–150.
- Engle, R. F. and Lee, G. (1999), *A long-run and short-run component model of stock return volatility*, Oxford University Press, pp. 475–497.
- Engle, R. F. and Ng, V. K. (1993), ‘Measuring and testing the impact of news on volatility’, *The Journal of Finance* **48**(5), 1749–1778.

- Engle, R., Ghysels, E. and Sohn, B. (2013), ‘Stock market volatility and macroeconomic fundamentals’, *Review of Economics and Statistics* **95**(3), 776–797.
- Engle, R. and Manganelli, S. (2004), ‘CAViaR: Conditional autoregressive value at risk by regression quantiles’, *Journal of Business and Economic Statistics* **22**(4), 367–381.
- Fantazzini, D. and Zimin, S. (2020), ‘A multivariate approach for the simultaneous modelling of market risk and credit risk for cryptocurrencies’, *Journal of Industrial and Business Economics* **47**(1), 19–69.
- Fasiolo, M. (2016), *An introduction to mvnfast. R package version 0.1.6*.
- Feign, A. (2021), ‘What is an ICO?’, CoinDesk. Available at: coindesk.com/learn/what-is-an-ico/.
- FINMA (2018), ‘Guidelines for enquiries regarding the regulatory framework for initial coin offerings (ICOs)’, Available at finma.ch/en/news/2018/02/20180216-mm-ico-wegleitung.
- Fisch, C. (2019), ‘Initial coin offerings (ICOs) to finance new ventures’, *Journal of Business Venturing* **34**(1), 1–22.
- Fisch, J. E., Gelbach, J. B. and Klick, J. (2017), ‘The logic and limits of event studies in securities fraud litigation’, *Tex. L. Rev.* **96**, 553.
- Fissler, T. and Ziegel, J. F. (2016), ‘Higher order elicibility and Osband’s principle’, *The Annals of Statistics* **44**(4), 1680–1707.
- Gaglianone, W. P., Lima, L. R., Linton, O. and Smith, D. R. (2011), ‘Evaluating value-at-risk models via quantile regression’, *Journal of Business & Economic Statistics* **29**(1), 150–160.
- Garcia-Bernardo, J., Fichtner, J., Takes, F. W. and Heemskerk, E. M. (2017), ‘Uncovering offshore financial centers: Conduits and sinks in the global corporate ownership network’, *Scientific Reports* **7**(1), 1–10.
- Gelbach, J. B., Helland, E. and Klick, J. (2013), ‘Valid inference in single-firm, single-event studies’, *American Law and Economics Review* **15**(2), 495–541.
- Ghalanos, A. (2019), *rmgarch: Multivariate GARCH models*. R package version 1.3-7.
- Ghalanos, A. (2020), *rugarch: Univariate GARCH models*. R package version 1.4-4.
- Giudici, G. and Adhami, S. (2019), ‘The impact of governance signals on ICO fundraising success’, *Journal of Industrial and Business Economics* **46**(2), 283–312.
- Glosten, L. R., Jagannathan, R. and Runkle, D. E. (1993), ‘On the relation between the expected value and the volatility of the nominal excess return on stocks’, *The Journal of Finance* **48**(5), 1779–1801.
- Gneiting, T. and Raftery, A. E. (2007), ‘Strictly proper scoring rules, prediction, and estimation’, *Journal of the American Statistical Association* **102**(477), 359–378.
- Gneiting, T. and Ranjan, R. (2011), ‘Comparing density forecasts using threshold- and quantile-weighted scoring rules’, *Journal of Business & Economic Statistics* **29**(3), 411–422.

- Griffin, J. M. and Shams, A. (2020), ‘Is bitcoin really untethered?’, *The Journal of Finance* **75**(4), 1913–1964.
- Grobys, K. and Huynh, T. (2021), ‘When tether says “JUMP!” bitcoin asks “How low?”’, *Finance Research Letters*.
- Guesmi, K., Saadi, S., Abid, I. and Ftiti, Z. (2019), ‘Portfolio diversification with virtual currency: Evidence from bitcoin’, *International Review of Financial Analysis* **63**, 431–437.
- Hair, J. F., ed. (2010), *Multivariate data analysis*, 7th edn, Prentice Hall.
- Hakala, S. (2017), ‘Lessons from single-company event studies: The importance of controlling for company-specific events’, ResearchGate. Working paper.
- Hampl, F. and Gyönyörová, L. (2021), ‘Can fiat-backed stablecoins be considered cash or cash equivalents under international financial reporting standards rules?’, *Australian Accounting Review* **31**(3), 233–255.
- Hansen, P., Lunde, A. and Nason, J. (2011), ‘The model confidence set’, *Econometrica* **79**(2), 453–497.
- Hansen, P. R., Huang, Z. and Shek, H. H. (2012), ‘Realized GARCH: A joint model for returns and realized measures of volatility’, *Journal of Applied Econometrics* **27**(6), 877–906.
- Hartmann, M. (2018), ‘This is how easy it is to buy ICO ratings - an investigation’. Available on Medium: <https://bit.ly/2IkUdVB>.
- Harvey, A. and Sucarrat, G. (2014), ‘EGARCH models with fat tails, skewness and leverage’, *Computational Statistics & Data Analysis* **76**, 320–338.
- Hattori, T. (2020), ‘A forecast comparison of volatility models using realized volatility: Evidence from the bitcoin market’, *Applied Economics Letters* **27**(7), 591–595.
- Hentschel, L. (1995), ‘All in the family: Nesting symmetric and asymmetric GARCH models’, *Journal of Financial Economics* **39**(1), 71–104.
- Hermans, L., Ianiro, A., Kochanska, U., Tormalehto, V.-M., van der Kraaij, A. and Simon, J. M. V. (2022), ‘Decrypting financial stability risks in crypto asset markets: Financial stability review, May 2022’, European Central Bank.
- Hertig, A. (2020), ‘What Is DeFi?’, CoinDesk.
- Higgins, M. L. and Bera, A. K. (1992), ‘A class of non-linear ARCH models’, *International Economic Review* **33**(1), 137–158.
- Hoang, L. and Baur, D. (2021), ‘How stable are stablecoins?’, *European Journal of Finance*. Article in Press.
- Hsu, D. and Ziedonis, R. (2013), ‘Resources as dual sources of advantage: Implications for valuing entrepreneurial-firm patents’, *Strategic Management Journal* **34**(7), 761–781.
- Huang, W., Meoli, M. and Vismara, S. (2020), ‘The geography of initial coin offerings’, *Small Business Economics* **55**(1), 77–102.

- Jordan, A., Krüger, F. and Lerch, S. (2019), ‘Evaluating probabilistic forecasts with scoringRules’, *Journal of Statistical Software* **90**(12), 1–37.
- Katsiampa, P. (2017), ‘Volatility estimation for bitcoin: A comparison of GARCH models’, *Economics Letters* **158**, 3–6.
- Köchling, G., Schmidtke, P. and Posch, P. (2020), ‘Volatility forecasting accuracy for bitcoin’, *Economics Letters* **191**, 108836.
- Kratz, M., Lok, Y. H. and McNeil, A. J. (2018), ‘Multinomial VaR backtests: A simple implicit approach to backtesting expected shortfall’, *Journal of Banking & Finance* **88**, 393–407.
- Kristoufek, L. (2021), ‘Tethered, or untethered? On the interplay between stablecoins and major cryptoassets’, *Finance Research Letters* **43**.
- Kristoufek, L. (2022), ‘On the role of stablecoins in cryptoasset pricing dynamics’, *Financial Innovation* **8**(1).
- Kupiec, P. (1995), ‘Techniques for verifying the accuracy of risk measurement models’, *Journal of Derivatives* **3**(2), 73–84.
- Lee, J., Li, T. and Shin, D. (2018), ‘The wisdom of crowds and information cascades in fintech: Evidence from initial coin offerings’. SSRN working paper 3226051.
- Leland, H. E. and Pyle, D. H. (1977), ‘Informational asymmetries, financial structure, and financial intermediation’, *The Journal of Finance* **32**(2), 371–387.
- Liu, W., Semeyutin, A., Lau, C. K. M. and Gozgor, G. (2020), ‘Forecasting Value-at-Risk of cryptocurrencies with RiskMetrics-type models’, *Research in International Business and Finance* **54**, 101259.
- Longerstaey, J. and Spencer, M. (1996), ‘RiskMetrics — Technical Document’, J. P. Morgan/Reuters.
- López-Cabarcos, M., Pérez-Pico, A., Piñeiro Chousa, J. and Šević, A. (2020), ‘Bitcoin volatility, stock market and investor sentiment. Are they connected?’, *Finance Research Letters* **38**, 101399.
- Lyandres, E., Palazzo, B. and Rabetti, D. (2022), ‘Initial coin offering (ICO) success and post-ICO performance’, *Management Science*. Article in Press.
- Lyons, R. K. and Viswanath-Natraj, G. (2020), ‘What keeps stablecoins stable?’, *NBER Working Paper* (w27136). Available at SSRN: <https://ssrn.com/abstract=3597868>.
- Maciel, L. (2021), ‘Cryptocurrencies value-at-risk and expected shortfall: Do regime-switching volatility models improve forecasting?’, *International Journal of Finance and Economics* **26**(3), 4840–4855.
- Maesa, D. D. F., Marino, A. and Ricci, L. (2017), ‘Detecting artificial behaviours in the bitcoin users graph’, *Online Social Networks and Media* **3-4**, 63 – 74.
- Matheson, J. E. and Winkler, R. L. (1976), ‘Scoring rules for continuous probability distributions’, *Management Science* **22**(10), 1087–1096.

- Matkovskyy, R., Jalan, A. and Dowling, M. (2020), 'Effects of economic policy uncertainty shocks on the interdependence between bitcoin and traditional financial markets', *Quarterly Review of Economics and Finance* **77**, 150–155.
- McAleer, M. and Da Veiga, B. (2008), 'Single-index and portfolio models for forecasting value-at-risk thresholds', *Journal of Forecasting* **27**(3), 217–235.
- McMillan, D. G. and Kambouroudis, D. (2009), 'Are RiskMetrics forecasts good enough? Evidence from 31 stock markets', *International Review of Financial Analysis* **18**(3), 117–124.
- McNeil, A. and Frey, R. (2000), 'Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach', *Journal of Empirical Finance* **7**(3–4), 271–300.
- Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M. and Savage, S. (2013), A fistful of bitcoins: Characterizing payments among men with no names, in 'Proceedings of the 2013 Conference on Internet Measurement Conference', Association for Computing Machinery, p. 127–140.
- Mincer, J. A. and Zarnowitz, V. (1969), The evaluation of economic forecasts, in 'Economic forecasts and expectations: Analysis of forecasting behavior and performance', NBER, pp. 3–46.
- Momtaz, P. (2021), 'Token offerings research database (TORD)'. Available at: <https://www.paulmomtaz.com/data/tord>.
- Momtaz, P. P. (2020a), 'Entrepreneurial finance and moral hazard: Evidence from token offerings', *Journal of Business Venturing* p. 106001. Article in press.
- Momtaz, P. P. (2020b), 'Initial coin offerings', *PLOS ONE* **15**(5), 1–30.
- Nakamoto, S. (2008), 'Bitcoin: A peer-to-peer electronic cash system'. Available at: bitcoin.org.
- Nekhili, R. and Sultan, J. (2020), 'Jump-driven risk model performance in cryptocurrency market', *International Journal of Financial Studies* **8**(2), 1–18.
- Nelson, D. B. (1991), 'Conditional heteroskedasticity in asset returns: A new approach', *Econometrica* **59**(2), 347–370.
- OECD (2019), 'Initial coin offerings (ICOs) for SME financing'. Report available at: <https://bit.ly/3pmY3yL>.
- Pafka, S. and Kondor, I. (2001), 'Evaluating the RiskMetrics methodology in measuring volatility and Value-at-Risk in financial markets', *Physica A: Statistical Mechanics and its Applications* **299**(1–2), 305–310.
- Patton, A. J. (2011), 'Volatility forecast comparison using imperfect volatility proxies', *Journal of Econometrics* **160**(1), 246–256.
- Philippi, S., Schuhmacher, M. C. and Bastian, N. (2021), 'Attracting investors in initial coin offerings: The relevance of specific technological capabilities for fundraising success', *Review of Corporate Finance* **1**(3–4), 455–485.

- Platt, C. (2017), ‘Thoughts on the taxonomy of blockchains & distributed ledger technologies’, Medium.
- Protos (2021a), ‘Tether papers: This is exactly who acquired 70% of all USDT ever issued’, Protos. Available at: bit.ly/3awMVg2.
- Protos (2021b), ‘Tether timeline: The complete history of crypto’s most stubborn stablecoin’, Protos. Available at: bit.ly/3OqLo9e.
- Puggioni, V. (2022), ‘Crypto rug pulls: What is a rug pull in crypto and 6 ways to spot it’, Cointelegraph.
- Rasskazova, A., Koroleva, E. and Rasskazov, S. (2019), Digital transformation: Statistical evaluation of success factors of an ICO campaign, in ‘2nd International Scientific Conference on Digital Transformation on Manufacturing, Infrastructure and Service, DTMIS 2018’, number 1 in ‘IOP Conference Series: Materials Science and Engineering’, IOP Publishing Ltd.
- Rhue, L. (2018), ‘Trust is all you need: An empirical exploration of initial coin offerings (ICOs) and ICO reputation scores’. SSRN working paper 3179723.
- Roosenboom, P., van der Kolk, T. and de Jong, A. (2020), ‘What determines success in initial coin offerings?’, *Venture Capital* **22**(2), 1–23.
- Rubinstein, R. Y. (1997), ‘Optimization of computer simulation models with rare events’, *European Journal of Operational Research* **99**(1), 89–112.
- Scheuerer, M. and Hamill, T. M. (2015), ‘Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities’, *Monthly Weather Review* **143**(4), 1321–1334.
- Schwert, G. W. (1990), ‘Stock volatility and the crash of ‘87’, *The Review of Financial Studies* **3**(1), 77–102.
- Segnon, M. and Bekiros, S. (2020), ‘Forecasting volatility in bitcoin market’, *Annals of Finance* **16**, 435–462.
- Sharma, Z. and Zhu, Y. (2020), ‘Platform building in initial coin offering market: Empirical evidence’, *Pacific-Basin Finance Journal* **61**, 101318.
- Shen, M. (2022), ‘How \$60 billion in terra coins went up in algorithmic smoke’, Bloomberg.
- Silahli, B., Dingec, K., Cifter, A. and Aydin, N. (2021), ‘Portfolio Value-at-Risk with two-sided Weibull distribution: Evidence from cryptocurrency markets’, *Finance Research Letters* **38**, 101425.
- Smith, S. (2021), ‘Decentralized finance and accounting – implications, considerations, and opportunities for development’, *International Journal of Digital Accounting Research* **21**, 129–153.
- Somin, S., Gordon, G. and Altshuler, Y. (2018), Network analysis of ERC20 tokens trading on Ethereum blockchain, in A. J. Morales, C. Gershenson, D. Braha, A. A. Minai and Y. Bar-Yam, eds, ‘Unifying Themes in Complex Systems IX’, Springer International Publishing, Cham, pp. 439–450.
- Sosa, M., Ortiz, E. and Cabello, A. (2019), ‘Bitcoin conditional volatility: GARCH extensions and Markov switching approach’, *International Finance Review* **20**, 201–219.

- Spence, M. (1973), 'Job market signaling', *The Quarterly Journal of Economics* **87**(3), 355–374.
- Spence, M. (2002), 'Signaling in retrospect and the informational structure of markets', *American Economic Review* **92**(3), 434–459.
- Stavroyiannis, S. (2018), 'Value-at-risk and related measures for the bitcoin', *The Journal of Risk Finance* **19**(2), 127–136.
- Szabo, N. (1997), 'Formalizing and securing relationships on public networks', *First Monday* **2**(9).
- Tasca, P., Hayes, A. and Liu, S. (2018), 'The evolution of the bitcoin economy: Extracting and analyzing the network of payment relationships', *The Journal of Risk Finance* **19**(2), 94 – 126.
- Thies, F., Wallbach, S., Wessel, M., Besler, M. and Benlian, A. (2021), 'Initial coin offerings and the cryptocurrency hype - the moderating role of exogenous and endogenous signals', *Electronic Markets* . Article in Press.
- Tiwari, A., Kumar, S. and Pathak, R. (2019), 'Modelling the dynamics of bitcoin and litecoin: GARCH versus stochastic volatility models', *Applied Economics* **51**(37), 4073–4082.
- Troster, V., Tiwari, A., Shahbaz, M. and Macedo, D. (2019), 'Bitcoin returns and risk: A general GARCH and GAS analysis', *Finance Research Letters* **30**, 187–193.
- Trucíos, C. (2019), 'Forecasting bitcoin risk measures: A robust approach', *International Journal of Forecasting* **35**(3), 836–847.
- Trucíos, C., Hotta, L. K. and Ruiz, E. (2017), 'Robust bootstrap forecast densities for GARCH returns and volatilities', *Journal of Statistical Computation and Simulation* **87**(16), 3152–3174.
- Trucíos, C. and Taylor, J. (2022), 'A comparison of methods for forecasting value-at-risk and expected shortfall of cryptocurrencies', ResearchGate Working Paper.
- Tse, Y. K. and Tsui, A. K. C. (2002), 'A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations', *Journal of Business & Economic Statistics* **20**(3), 351–362.
- Twomey, D. and Mann, A. (2020), Fraud and manipulation within cryptocurrency markets, in C. Alexander and D. Cumming, eds, 'Fraud and Corruption in Financial Markets'.
- Vicati, A. (2018), 'A closer look at tether's blockchain', Hacker Noon. Available at: bit.ly/3aDRftI.
- Vidal-Tomás, D. and Ibañez, A. (2018), 'Semi-strong efficiency of bitcoin', *Finance Research Letters* **27**, 259–265.
- Wang, G.-J., Ma, X.-Y. and Wu, H.-Y. (2020), 'Are stablecoins truly diversifiers, hedges, or safe havens against traditional cryptocurrencies as their name suggests?', *Research in International Business and Finance* **54**.
- Wang, J.-N., Liu, H.-C., Chiang, S.-M. and Hsu, Y.-T. (2019), 'On the predictive power of ARJI volatility forecasts for bitcoin', *Applied Economics* **51**(44), 4849–4855.

- Wei, W. (2018), ‘The impact of tether grants on bitcoin’, *Economics Letters* **171**, 19–22.
- Weyl, E. G., Ohlhaver, P. and Buterin, V. (2022), ‘Decentralized society: Finding web3’s soul’, SSRN Discussion Paper 4105763.
- Wood, G. (2014), ‘Ethereum: A secure decentralised generalised transaction ledger’, GitHub, v151.2014. Available at: github.com/ethereum/yellowpaper.
- Yen, J.-C., Wang, T. and Chen, Y.-H. (2021), ‘Different is better: How unique initial coin offering language in white papers enhances success’, *Accounting and Finance* **61**(4), 5309–5340.
- Zakoian, J.-M. (1994), ‘Threshold heteroskedastic models’, *Journal of Economic Dynamics and Control* **18**(5), 931–955.
- Zoromé, A. (2007), ‘Concept of offshore financial centers: In search of an operational definition’. IMF working paper 07/87.

SUPPLEMENTARY MATERIAL – CRYPTO ASSET MARKET RISK MODELLING

Table A1: Further backtesting of daily left-tail 1% VaR for a short position

	BTC	ETH	XRP	LTC
$t\text{EWMA}(0.925, 6)$	0.9962	0.9923	1.0000	1.0000
$t\text{EWMA}(0.94, 6)$	0.9962	0.9962	1.0000	0.9999
$t\text{AEWMA}(0.925, 1\%, 6)$	0.9992	0.9923	1.0000	0.9999
$t\text{AEWMA}(0.925, 2\%, 6)$	0.9549	0.9737	1.0000	0.9999
$t\text{AEWMA}(0.925, 3\%, 6)$	0.5924	0.8864	1.0000	0.9854
$t\text{AEWMA}(0.925, 4\%, 6)$	0.0707	0.5924	0.9992	0.8325
$t\text{AEWMA}(0.925, 5\%, 6)$	0.0036	0.1627	0.9549	0.2296
$t\text{AEWMA}(0.94, 1\%, 6)$	0.9962	0.9854	1.0000	0.9999
$t\text{AEWMA}(0.94, 2\%, 6)$	0.9266	0.9266	1.0000	0.9982
$t\text{AEWMA}(0.94, 3\%, 6)$	0.2296	0.8325	1.0000	0.9854
$t\text{AEWMA}(0.94, 4\%, 6)$	0.0251	0.5924	0.9992	0.6834
$t\text{AEWMA}(0.94, 5\%, 6)$	0.0036	0.1099	0.9737	0.3097
$t\text{AEWMA}(0.925, -1\%, 6)$	0.8325	0.9737	1.0000	1.0000
$t\text{AEWMA}(0.925, -2\%, 6)$	0.4961	0.7644	1.0000	0.9923
$t\text{AEWMA}(0.925, -3\%, 6)$	0.1627	0.2296	0.9982	0.9549
$t\text{AEWMA}(0.925, -4\%, 6)$	0.0073	0.0139	0.8864	0.4001
$t\text{AEWMA}(0.925, -5\%, 6)$	0.0017	0.0007	0.4961	0.0707
$t\text{AEWMA}(0.94, -1\%, 6)$	0.7644	0.9854	1.0000	0.9999
$t\text{AEWMA}(0.94, -2\%, 6)$	0.4001	0.7644	1.0000	0.9737
$t\text{AEWMA}(0.94, -3\%, 6)$	0.1099	0.1627	0.9982	0.8864
$t\text{AEWMA}(0.94, -4\%, 6)$	0.0073	0.0073	0.9266	0.4001
$t\text{AEWMA}(0.94, -5\%, 6)$	0.0007	0.0007	0.5924	0.0433

Backtesting of daily left-tail 1% VaR for a short position on bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC). The table reports for each asset the VaR traffic light backtest probability $\Phi(z)$ of obtaining the observed number of exceedances or fewer for the 1% right-tail VaR for various EWMA and AEWMA model specifications with the Student- t degrees of freedom parameter set ad hoc to $\nu = 6$.

Table A2: Further backtesting of hourly left-tail 1% VaR

	BTC	ETH	XRP	LTC
$tEWMA(0.925)$	0.9999	0.9999	0.9999	0.9986
$tEWMA(0.94)$	0.9998	0.9994	0.9994	0.9930
$tAEWMA(0.925, 0.2\%)$	0.9998	0.9986	0.9994	0.9930
$tAEWMA(0.925, 0.4\%)$	0.9858	0.8733	0.9986	0.8733
$tAEWMA(0.925, 0.6\%)$	0.7314	0.6385	0.9732	0.6385
$tAEWMA(0.925, 0.7\%)$	0.5366	0.5366	0.9523	0.4322
$tAEWMA(0.925, 0.8\%)$	0.1689	0.4322	0.8733	0.3324
$tAEWMA(0.925, 0.9\%)$	0.0690	0.3324	0.8105	0.3324
$tAEWMA(0.925, 1\%)$	0.0404	0.1689	0.6385	0.1689
$tAEWMA(0.925, 1.2\%)$	0.0056	0.0404	0.5366	0.0690
$tAEWMA(0.925, 1.4\%)$	0.0026	0.0116	0.3324	0.0223
$tAEWMA(0.925, 1.6\%)$	0.0011	0.0056	0.1110	0.0223
$tAEWMA(0.925, 1.8\%)$	0.0011	0.0011	0.0223	0.0116
$tAEWMA(0.925, 2\%)$	0.0011	0.0011	0.0223	0.0011
$tAEWMA(0.925, -0.2\%)$	1.0000	0.9994	1.0000	0.9986
$tAEWMA(0.925, -0.4\%)$	0.9967	0.9858	0.9998	0.9930
$tAEWMA(0.925, -0.6\%)$	0.8733	0.8733	0.9930	0.9732
$tAEWMA(0.925, -0.7\%)$	0.7314	0.8733	0.9858	0.9732
$tAEWMA(0.925, -0.8\%)$	0.7314	0.6385	0.9732	0.9732
$tAEWMA(0.925, -0.9\%)$	0.5366	0.5366	0.8105	0.9523
$tAEWMA(0.925, -1\%)$	0.4322	0.3324	0.5366	0.6385
$tAEWMA(0.94, 0.2\%)$	0.9994	0.9930	0.9967	0.9858
$tAEWMA(0.94, 0.4\%)$	0.9732	0.8105	0.9967	0.9200
$tAEWMA(0.94, 0.6\%)$	0.6385	0.5366	0.9732	0.6385
$tAEWMA(0.94, 0.7\%)$	0.5366	0.5366	0.9200	0.4322
$tAEWMA(0.94, 0.8\%)$	0.0690	0.4322	0.8105	0.4322
$tAEWMA(0.94, 0.9\%)$	0.0404	0.1689	0.8105	0.2433
$tAEWMA(0.94, 1\%)$	0.0223	0.0690	0.6385	0.2433
$tAEWMA(0.94, 1.2\%)$	0.0056	0.0223	0.4322	0.1110
$tAEWMA(0.94, 1.4\%)$	0.0026	0.0056	0.2433	0.0223
$tAEWMA(0.94, 1.6\%)$	0.0011	0.0056	0.1110	0.0223
$tAEWMA(0.94, 1.8\%)$	0.0011	0.0026	0.0223	0.0056
$tAEWMA(0.94, 2\%)$	0.0011	0.0011	0.0223	0.0011
$tAEWMA(0.94, -0.2\%)$	0.9998	0.9986	0.9994	0.9930
$tAEWMA(0.94, -0.4\%)$	0.9858	0.9732	0.9994	0.9930
$tAEWMA(0.94, -0.6\%)$	0.8105	0.7314	0.9967	0.9732
$tAEWMA(0.94, -0.7\%)$	0.7314	0.6385	0.9732	0.9732
$tAEWMA(0.94, -0.8\%)$	0.7314	0.4322	0.9200	0.9523
$tAEWMA(0.94, -0.9\%)$	0.5366	0.3324	0.5366	0.8105
$tAEWMA(0.94, -1\%)$	0.4322	0.1689	0.4322	0.7314

Backtesting of hourly left-tail 1% VaR on bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC). The table reports for each asset the VaR traffic light backtest probability $\Phi(z)$ of obtaining the observed number of exceedances or fewer for the 1% left-tail VaR for various EWMA and AEWMA model specifications with the Student- t degrees of freedom parameter set ad hoc to $\nu = 6$.

Table A3: Further backtesting of hourly right-tail 1% VaR

	BTC	ETH	XRP	LTC
t EWMA(0.925)	0.6385	0.6385	0.9200	0.6385
t EWMA(0.94)	0.5366	0.5366	0.7314	0.6385
t AEWMA(0.925, 0.2%)	0.7314	0.7314	0.7314	0.6385
t AEWMA(0.925, 0.4%)	0.2433	0.5366	0.4322	0.6385
t AEWMA(0.925, 0.6%)	0.0223	0.0690	0.1689	0.2433
t AEWMA(0.925, 0.7%)	0.0223	0.0404	0.1110	0.1689
t AEWMA(0.925, 0.8%)	0.0116	0.0404	0.1110	0.1110
t AEWMA(0.925, 0.9%)	0.0056	0.0404	0.1110	0.0690
t AEWMA(0.925, 1%)	0.0026	0.0116	0.0690	0.0116
t AEWMA(0.925, 1.2%)	0.0011	0.0056	0.0223	0.0056
t AEWMA(0.925, 1.4%)	0.0011	0.0011	0.0056	0.0056
t AEWMA(0.925, 1.6%)	0.0004	0.0004	0.0056	0.0011
t AEWMA(0.925, 1.8%)	0.0002	0.0002	0.0026	0.0002
t AEWMA(0.925, 2%)	0.0001	0.0002	0.0004	0.0002
t AEWMA(0.925, -0.2%)	0.5366	0.6385	0.8105	0.7314
t AEWMA(0.925, -0.4%)	0.4322	0.5366	0.6385	0.3324
t AEWMA(0.925, -0.6%)	0.1110	0.1689	0.2433	0.0690
t AEWMA(0.925, -0.7%)	0.0690	0.1110	0.2433	0.0404
t AEWMA(0.925, -0.8%)	0.0116	0.0690	0.1110	0.0404
t AEWMA(0.925, -0.9%)	0.0056	0.0690	0.1110	0.0404
t AEWMA(0.925, -1%)	0.0026	0.0116	0.1110	0.0404
t AEWMA(0.94, 0.2%)	0.2433	0.7314	0.7314	0.6385
t AEWMA(0.94, 0.4%)	0.1689	0.4322	0.3324	0.5366
t AEWMA(0.94, 0.6%)	0.0223	0.0690	0.1110	0.3324
t AEWMA(0.94, 0.7%)	0.0116	0.0404	0.1110	0.2433
t AEWMA(0.94, 0.8%)	0.0026	0.0404	0.1110	0.1110
t AEWMA(0.94, 0.9%)	0.0026	0.0404	0.0690	0.0223
t AEWMA(0.94, 1%)	0.0026	0.0116	0.0404	0.0223
t AEWMA(0.94, 1.2%)	0.0011	0.0056	0.0223	0.0056
t AEWMA(0.94, 1.4%)	0.0011	0.0011	0.0056	0.0026
t AEWMA(0.94, 1.6%)	0.0004	0.0004	0.0056	0.0004
t AEWMA(0.94, 1.8%)	0.0001	0.0002	0.0026	0.0002
t AEWMA(0.94, 2%)	0.0001	0.0002	0.0004	0.0002
t AEWMA(0.94, -0.2%)	0.5366	0.5366	0.6385	0.6385
t AEWMA(0.94, -0.4%)	0.3324	0.3324	0.5366	0.2433
t AEWMA(0.94, -0.6%)	0.0690	0.1689	0.1689	0.1110
t AEWMA(0.94, -0.7%)	0.0404	0.1110	0.1110	0.0404
t AEWMA(0.94, -0.8%)	0.0116	0.0404	0.1110	0.0404
t AEWMA(0.94, -0.9%)	0.0116	0.0223	0.1110	0.0404
t AEWMA(0.94, -1%)	0.0056	0.0116	0.1110	0.0223

Backtesting of hourly right-tail 1% VaR on bitcoin (BTC), ether (ETH), ripple (XRP) and litecoin (LTC). The table reports for each asset the VaR traffic light backtest probability $\Phi(z)$ of obtaining the observed number of exceedances or fewer for the 1% right-tail VaR for various EWMA and AEWMA model specifications with the Student- t degrees of freedom parameter set ad hoc to $\nu = 6$.


SUPPLEMENTARY MATERIAL – BLOCKCHAIN TRANSACTIONS AND CRYPTO ASSET MARKET MANIPULATION

Table B1: Material USDT outflow events from the tether treasury towards exchanges (MVDA)

Start	End	USDT	$\hat{\gamma}$	SQ	t_{γ}	SQ_{std}	Adj. R^2	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-09-24 20:00:00	2019-09-24 22:00:00	9,990,000	0.005	0.002	4.72	1.48	0.93	8,595.3	106.8	39.4	0.4	1,257,490
2020-03-16 17:00:00	2020-03-16 19:00:00	6,708,292	0.011	0.004	4.63	0.82	0.94	5,089.5	-134.9	58.2	-0.4	974,281
2020-03-17 14:00:00	2020-03-17 16:00:00	8,841,159	0.011	0.004	3.55	0.71	0.91	5,254.1	121.1	57.7	0.5	1,784,288
2020-07-26 09:00:00	2020-07-26 11:00:00	10,100,000	0.002	0.001	4.42	1.43	0.93	9,846.6	132.0	17.7	0.1	183,582
2020-07-27 05:00:00	2020-07-27 07:00:00	12,000,000	0.004	0.001	8.60	1.12	0.90	10,299.7	-129.0	46.2	-0.4	377,491
2020-07-27 21:00:00	2020-07-27 23:00:00	35,000,000	0.003	0.001	4.41	0.98	0.87	11,120.3	-77.4	34.6	-0.4	425,854
2020-08-01 06:00:00	2020-08-01 08:00:00	15,000,000	0.002	0.002	2.49	1.18	0.89	11,648.1	34.3	23.0	0.7	312,540
2020-09-15 12:00:00	2020-09-15 14:00:00	19,000,000	0.002	0.002	2.63	1.31	0.92	10,914.9	-141.4	24.0	-0.2	194,605
2020-10-19 15:00:00	2020-10-19 17:00:00	10,000,000	0.003	0.001	5.30	1.51	0.95	11,804.9	-75.4	29.8	-0.4	322,961
2020-10-27 15:00:00	2020-10-27 17:00:00	15,000,000	0.002	0.002	3.09	1.51	0.94	13,618.0	82.3	25.0	0.3	138,237
2020-11-30 14:00:00	2020-11-30 16:00:00	26,000,000	0.004	0.003	2.64	1.24	0.87	19,494.6	-304.6	77.0	-0.3	1,935,313

Material USDT outflow events from the tether treasury towards exchanges and corresponding estimates of the cap-weighted MVDA index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_{γ} , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.

Figure B1: Hourly returns (PCA)



Hourly arithmetic returns on the first 4 principal components, based on the returns on the USD-denominated prices of 135 crypto assets obtained from Cryptocompare. The sample period is 3 March 2019 - 1 January 2021.

Table B2: Material USDT outflow events from the tether treasury towards exchanges (PCA)

Start	End	USDT	$\hat{\gamma}$	SQ	t_γ	SQ_{std}	Adj. R^2	V_{PCA}	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-09-19 21:00:00	2019-09-19 23:00:00	4,995,000	0.007	0.007	2.37	1.33	0.01	0.51	10,277.6	3.5	76.1	21.9	343,632
2019-09-25 21:00:00	2019-09-25 23:00:00	5,000,000	0.010	0.007	2.99	1.18	0.05	0.45	8,602.1	-155.1	86.3	-0.6	551,071
2020-07-23 16:00:00	2020-07-23 18:00:00	23,000,000	0.004	0.004	2.83	1.32	0.16	0.53	9,631.6	-41.0	42.8	-1.0	391,880
2020-07-26 09:00:00	2020-07-26 11:00:00	10,100,000	0.005	0.003	3.70	1.32	0.10	0.51	9,846.6	132.0	53.0	0.4	548,080
2020-07-27 05:00:00	2020-07-27 07:00:00	12,000,000	0.008	0.004	5.36	1.40	0.10	0.54	10,299.7	-129.0	87.4	-0.7	714,884
2020-07-27 21:00:00	2020-07-27 23:00:00	35,000,000	0.008	0.004	4.42	1.23	0.09	0.55	11,120.3	-77.4	93.7	-1.2	1,151,813
2020-07-28 15:00:00	2020-07-28 19:00:00	27,000,000	0.006	0.005	3.79	1.26	0.05	0.55	11,172.8	-160.8	68.8	-0.4	525,536
2020-08-01 06:00:00	2020-08-01 08:00:00	15,000,000	0.008	0.005	3.56	1.35	0.07	0.52	11,648.1	34.3	95.1	2.8	1,294,061
2020-09-15 12:00:00	2020-09-15 14:00:00	19,000,000	0.005	0.004	3.00	1.35	0.65	0.46	10,914.9	-141.4	58.6	-0.4	476,152
2020-10-19 15:00:00	2020-10-19 17:00:00	10,000,000	0.007	0.003	5.08	1.24	0.56	0.38	11,804.9	-75.4	82.0	-1.1	890,514
2020-10-21 12:00:00	2020-10-21 16:00:00	29,000,000	0.003	0.003	2.97	1.31	0.51	0.38	12,423.1	417.5	41.9	0.1	568,026
2020-10-27 15:00:00	2020-10-27 17:00:00	15,000,000	0.005	0.004	2.53	1.26	0.36	0.38	13,618.0	82.3	66.5	0.8	368,121
2020-10-29 15:00:00	2020-10-29 17:00:00	15,000,000	0.006	0.005	2.83	1.29	0.35	0.37	13,440.6	148.8	77.5	0.5	528,110
2020-11-18 08:00:00	2020-11-18 10:00:00	14,500,000	0.011	0.009	3.17	1.54	0.04	0.42	18,219.1	33.8	200.9	6.0	1,898,746
2020-11-30 14:00:00	2020-11-30 16:00:00	26,000,000	0.016	0.010	4.53	1.57	0.22	0.38	19,494.6	-304.6	321.2	-1.1	8,075,951
2020-12-19 14:00:00	2020-12-19 16:00:00	27,000,000	0.008	0.007	3.03	1.55	0.57	0.50	23,304.2	679.2	195.5	0.3	2,848,536

Material USDT outflow events from the tether treasury towards exchanges and corresponding estimates of the PCA-based 4-factor index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_γ , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.

Table B3: Material USDT outflow events from the tether treasury towards unknown entities (MVDA)

Start	End	USDT	$\hat{\gamma}$	SQ	t_γ	SQ_{std}	Adj. R^2	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-05-13 17:00:00	2019-05-13 19:00:00	4,995,000	0.005	0.004	3.07	1.55	0.81	7,800.1	82.7	35.7	0.4	581,025
2020-03-17 16:00:00	2020-03-17 23:00:00	83,500,000	0.006	0.004	2.93	0.67	0.91	5,375.3	-37.6	30.0	-0.8	1,035,619
2020-03-20 02:00:00	2020-03-20 04:00:00	15,000,000	0.010	0.004	2.32	0.58	0.83	6,159.2	41.6	62.0	1.5	642,195
2020-10-21 14:00:00	2020-10-21 16:00:00	10,000,000	0.003	0.001	5.05	1.53	0.94	12,741.4	99.1	34.3	0.3	598,746
2020-10-21 22:00:00	2020-10-22 00:00:00	25,000,000	0.004	0.001	7.78	1.48	0.94	13,219.7	-298.0	56.8	-0.2	846,659
2020-11-30 14:00:00	2020-11-30 16:00:00	49,999,999	0.004	0.003	2.59	1.28	0.87	19,494.6	-304.6	75.4	-0.2	1,896,535
2020-12-19 17:00:00	2020-12-19 19:00:00	100,000,000	0.003	0.003	2.01	1.11	0.88	23,908.0	-96.4	69.7	-0.7	586,547

Material USDT outflow events from the tether treasury towards other unknown entities and corresponding estimates of the cap-weighted MVDA index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_γ , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the ratio of total system variance V_{PCA} explained by the 4 principal components, the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.


Table B4: Material USDT outflow events from the tether treasury towards unknown entities (PCA)

Start	End	USDT	$\hat{\gamma}$	SQ	t_{γ}	SQ_{std}	Adj. R^2	V_{PCA}	P_{start}	ΔP	P^{Effect}	Ratio	Damage
2019-05-13 17:00:00	2019-05-13 19:00:00	4,995,000	0.015	0.008	4.82	1.55	0.22	0.33	7,800.1	82.7	113.4	1.4	1,847,047
2020-04-23 14:00:00	2020-04-23 17:00:00	25,584,859	0.016	0.011	4.41	1.46	0.07	0.40	7,548.3	11.7	120.1	10.2	7,057,017
2020-04-29 20:00:00	2020-04-30 00:00:00	22,001,000	0.013	0.009	4.44	1.44	0.10	0.44	8,931.3	-115.1	112.7	-1.0	1,928,239
2020-05-07 17:00:00	2020-05-07 22:00:00	36,000,000	0.010	0.010	3.52	1.43	0.02	0.73	9,844.1	55.5	97.2	1.8	4,769,304
2020-06-01 23:00:00	2020-06-02 03:00:00	50,000,000	0.013	0.010	3.77	1.34	0.07	0.75	10,211.2	-116.3	131.6	-1.1	5,965,512
2020-07-22 22:00:00	2020-07-23 00:00:00	15,000,000	0.005	0.004	3.14	1.41	0.16	0.54	9,507.4	17.0	46.7	2.7	547,138
2020-10-21 14:00:00	2020-10-21 16:00:00	10,000,000	0.009	0.003	6.03	1.25	0.53	0.38	12,741.4	99.1	112.6	1.1	1,964,601
2020-10-21 22:00:00	2020-10-22 00:00:00	25,000,000	0.012	0.003	7.94	1.32	0.50	0.38	13,219.7	-298.0	161.3	-0.5	2,405,791
2020-11-30 14:00:00	2020-11-30 16:00:00	49,999,999	0.016	0.010	4.46	1.59	0.22	0.38	19,494.6	-304.6	316.6	-1.0	7,960,818

Material USDT outflow events from the tether treasury towards other unknown entities and corresponding estimates of the PCA-based 4-factor index model, for which the alternate hypothesis of a significant positive abnormal return is accepted at the 5% significance level simultaneously for the parametric t-test – using a critical value on the standardized Student- t distribution with the degrees of freedom parameter determined via the regression model, and the raw and standardized SQ tests. The columns denote: the start and end of each USDT outflow event, the amount of USDT sent from the tether treasury, the estimated average abnormal return $\hat{\gamma}$ during the event, the raw SQ test critical value based on the 5% right-tail quantile of the estimated regression residuals, the estimated t-ratio of the average abnormal return t_{γ} , the standardized SQ test critical value based on the estimated residuals scaled by their estimated standard error, the regression adjusted R^2 , the ratio of total system variance V_{PCA} explained by the 4 principal components, the BTC/USD price P_{start} at the beginning of each event, the change in price ΔP between the end and beginning of each event, the estimated price effect P^{Effect} of each event, the ratio $P^{Effect}/\Delta P$ and the estimated damage calculated as the product of the price effect and total traded volume at the beginning of the event, divided by 2.


SUPPLEMENTARY MATERIAL – THE TOKENOMICS OF CROWDFUNDING

Figure C1: Token offering launchpad platform coefficients




Constant coefficients for the token offering platform based on the linear regression model with $\log\text{Raised}$ as the dependent variable. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.

Figure C2: Token offerings per launchpad platform



Number of token offerings per platform. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.

Figure C3: Token offering category coefficients



Constant coefficients for the token offering category based on the linear regression model with logRaised as the dependent variable. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.

Table C1: Linear regression results for the 2017 – 2020 period.

	(i) Basic		(ii) Full	
Constant	-0.696***	(-4.42)	-0.584***	(-3.02)
logtCap	0.332***	(9.33)	0.327***	(9.40)
Presale	-0.061	(-1.18)	-0.070	(-1.33)
Rating	0.196***	(4.51)	0.190***	(4.32)
Team	0.021***	(6.00)	0.021***	(5.71)
Advisors	0.012***	(2.73)	0.011**	(2.49)
Bonus	-0.110**	(-2.15)	-0.124**	(-2.34)
Distribution	0.076**	(2.53)	0.062**	(2.00)
logDuration	-0.060***	(-3.01)	-0.089***	(-4.00)
ETH price	0.182***	(7.63)	0.171***	(6.58)
ETH mom	-0.041	(-1.60)	-0.023	(-0.87)
OFC	0.187***	(3.73)	0.168***	(3.29)
Accepts BTC			0.208***	(3.63)
Accepts ETH			-0.018	(-0.18)
Accepts other			-0.088	(-1.51)
Accepts fiat			-0.063	(-0.88)
Capped			-0.013	(-0.14)
Ethereum-based			0.023	(0.30)
IEO			-0.248***	(-2.61)
KYC			0.063	(1.03)
Whitelist			0.033	(0.61)
High			0.004	(0.08)
Medium			-0.061	(-1.20)
Low			-0.035	(-0.46)
Observations	1319		1319	
R^2	0.232		0.245	
Adj. R^2	0.226		0.232	

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Linear regression results with logRaised as the dependent variable. The model specifications are: (i) Basic which only includes the variables involved in our hypotheses; (ii) Full which further includes control variables. The first column in each model specification displays the regression betas and the second column displays the t-statistics in parentheses. A sample of 1,319 ICOs is used and all ICOs take place between 1 January 2017 - 31 December 2020.


Table C2: Probit regression results for the 2017 – 2020 period.

	(i) Basic		(ii) Full	
logSoftcap	-0.131***	(-6.53)	-0.141***	(-6.81)
logtCap	0.077***	(3.65)	0.075***	(3.52)
Presale (d)	-0.018	(-0.53)	-0.024	(-0.68)
Rating	0.106***	(3.62)	0.106***	(3.44)
Team	0.007**	(2.50)	0.007**	(2.47)
Advisors	0.004	(1.46)	0.005	(1.51)
ETH price	0.030	(1.61)	0.023	(1.17)
ETH mom	-0.011	(-0.63)	-0.001	(-0.08)
OFC (d)	0.093***	(2.61)	0.089**	(2.46)
Bonus (d)	-0.076**	(-2.18)	-0.073**	(-2.02)
Distribution	-0.015	(-0.82)	-0.022	(-1.16)
logDuration	-0.036**	(-2.23)	-0.057***	(-3.24)
Accepts BTC (d)			0.137***	(3.30)
Accepts ETH (d)			-0.088	(-1.32)
Accepts other (d)			-0.028	(-0.66)
Accepts fiat (d)			0.002	(0.04)
Ethereum-based (d)			0.056	(1.14)
IEO (d)			-0.202***	(-2.60)
KYC (d)			0.031	(0.75)
Whitelist (d)			-0.026	(-0.64)
High (d)			-0.010	(-0.27)
Medium (d)			-0.038	(-1.09)
Low (d)			0.001	(0.02)
Observations	902		902	
Pseudo- R^2	0.086		0.104	

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Probit regression marginal effects at means (MEM) with soft cap exceedance as the dependent variable. The model specifications are: (i) Basic which only includes the variables involved in our hypotheses; (ii) Full which further includes control variables. The notation (d) next to a variable indicates that the corresponding marginal effect is calculated for a discrete change of the binary variable from 0 to 1. The first column in each model specification displays the marginal effect and the second column displays the z-statistics in parentheses. A sample of 902 ICOs is used and all ICOs take place between 1 January 2017 - 31 December 2020.

Figure C4: Token offerings per category



Number of token offerings per category. The sample period is 1 January 2021 - 31 January 2022 containing 1,607 offerings.