Thema:

Valuation Models for Crypto Assets:
The Case of Cryptocurrencies

(Bewertungsmodelle für Krypto-Assets
am Beispiel von Kryptowährungen)

Masterarbeit

im Rahmen des CUR Executive Accounting & Controlling Program

Themensteller: Prof. Dr. Wolfgang Berens
Betreuer: Simon Schölzel, M.Sc.

vorgelegt von: Dr. Benjamin Jurke
Frieda-Nadig-Str. 45
33332 Gütersloh
+49-151-57359380
mail@benjaminjurke.com

Abgabetermin: 2022-08-15

Table of Contents

List of Abbreviations

AMM	automated market maker
ARIMA	auto-regressive integrated moving average
APY	annual percentage yield
CBOE	Chicago Board Options Exchange
CME	Chicago Mercantile Exchange
CNN	convolutional neural network
DAO	decentralized autonomous organization
DeFi	decentralized finance
DEX	decentralized exchange
FBI	Federal Bureau of Investigation
LSTM	long-short-term memory (neural network)
NFT	non-fungible token
NN	neural network
NVML	network value to Metcalfe's law
NVOL	network value to Odlyzko's law
NVT	network value to transactions
RNN	recurrent neural network
SARIMA	seasonal auto-regressive integrated moving average
SF / S2F	stock-to-flow
TLCC	time-lagged cross-correlation
TVL	total value locked (DeFi metric)
Tx	transaction(s)

List of Symbols

BTC	Bitcoin / Bitcoin blockchain token
BTCUSDT	Bitcoin-Tether futures contract on Binance exchange
DAI	USD-pegged stablecoin of MakerDAO project
ETH	Ether / Ethereum blockchain token
MATIC	Polygon blockchain token
MKR	governance token of MakerDAO project
SBNY	Signature Bank of New York trading symbol
SI	Silvergate Bank trading symbol
USD	US dollar
USDC	USD-pegged stablecoin of Binance exchange
USDT	USD-pegged stablecoin of Tether
WBTC	wrapped BTC

List of Figures

All graphs have been created by the author using the Python package “matplotlib”.

List of Tables

1 Introduction

In late October 2008 a whitepaper[1] appeared on the "Cryptographic Mailing List", which laid out the technological foundation of blockchain technology and introduced the Bitcoin cryptocurrency to the world. The paper has been referred to as "the most important contribution to monetary economics in the 21st century"[2]. Just one month earlier the fourth-largest investment bank Lehman Brothers had collapsed, ushering in the biggest financial crisis in recent decades.

Using a combination of cryptographic elements like hashes and signatures together with a novel synchronization mechanism based on solving computationally expensive puzzles, the introduction of Bitcoin marks the beginning of the first decentralized transaction system for a digital currency. The blockchain technology removes the necessity for an intermediary third-party trustee between two transacting participants, like a bank, and is therefore often regarded as the first trustless peer-to-peer digital payment system. More generally, this allows to reach a consensus between potentially distrusting parties about the ownership of a digital or virtual asset without the need for centralization. From the technological point of view a blockchain represents a decentralized append-only database which solves the double spending problem, i.e., it ensures that a digital asset can only be spent once during a transaction, thus introducing scarcity for virtual goods.

In the history of money, the emergence of a monetary unit appears to be a spontaneous process when a certain good in a previously barter-based economy is traded frequently[3]. By using a generally accepted medium of exchange, the number of potential trading pairs grows linearly instead of quadratically with an increasing number of produced goods, which makes complex economies feasible. It also allows a decoupling from the simultaneous occurrences of mutual needs, especially when perishable goods are exchanged. Furthermore, aside from being a medium of exchange and a unit of account, a monetary unit also provides a store of value between the time it is received and the time it is spend again.

[1] See reference (Nakamoto 2008).
[2] See introduction in (Schär and Berentsen 2020).
[3] See reference (Menger 1892).

The market value of a monetary unit can be separated into three components: the intrinsic value of the material in case of a physically exchanged medium, the value of an attached promise of payment and the liquidity premium, which represents the expected future acceptance of the monetary unit in exchange for goods and services[4]. By overturning the Bretton-Woods convertibility of US dollars into gold in 1971, the United States turned the US dollar and all of the attached major currencies of the world into fiat money, which has no intrinsic value and offers no promise of payment. While the naturally limited quantities and yearly increments produced by the mining of gold put a boundary on the number of monetary units issued by the central bank, the untied fiat currencies led to an expansion of the global debt level approaching post-WW2 levels since the 1970s as measured in terms of the worlds global economic output, especially in the developed nations. Various subsequent deregulations in the financial industry allowed banks to engage in hedge fund-like trading activities and repackaging of presumably low-risk mortgages to increase their leverage and profitability, which led to the creation of interest-only loans affordable to subprime borrowers without risk-adjusted premiums. The 2008 financial crisis, following the collapse of the housing bubble created by the unchecked lending practices and the subsequent liquidity provisioning of the central banks to limit the economic turmoil, set the stage in which the publication of the Nakamoto paper occurred.

The Bitcoin blockchain is set up with a limited supply of 21 million BTC, the symbol for a single unit of the virtual currency. Since all of those will eventually be mined in a gradually slowing process over time, a Bitcoin-based monetary system exhibits both a decreasing inflation in the overall supply of monetary units as well as deflationary aspects due to the decreasing amounts of newly created BTC units over time. Therefore, Bitcoin is often referred to as digital gold due to the perceived value embodied by the artificially imposed scarcity and increasing hardness to obtain further units.

Following the release of numerous copycat projects to the Bitcoin ecosystem, the first truly novel extension was introduced with the Ethereum whitepaper[5]. While

---

[4] See section 1.4 in (Schär and Berentsen 2020).
[5] See reference (Buterin 2014).

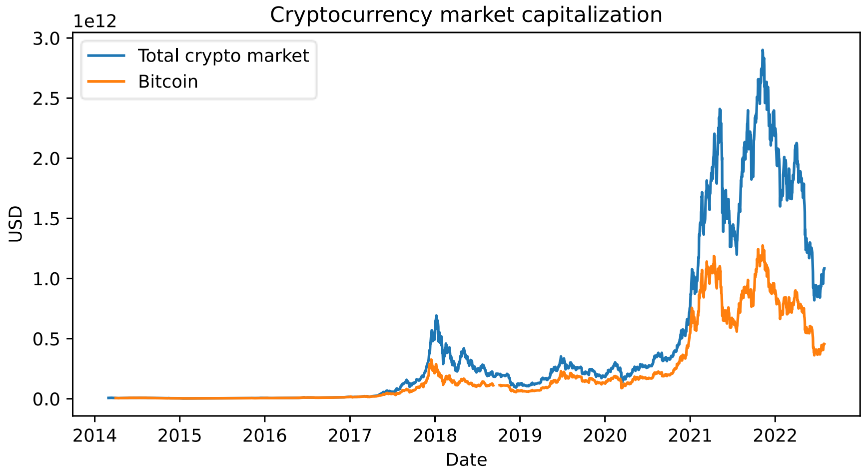

Figure 1: Total cryptocurrency and Bitcoin market capitalization
(Created using data retrieved from TradingView, see https://www.tradingview.com.)

the Bitcoin network protocol is restrictive, purpose-build and essentially only allows for the transactional character of fungible tokens to take place, Ethereum introduced the concept of distributed general-purpose computation and smart contracts to the blockchain ecosystem. This in turn led to the development of non-fungible tokens, which represent uniquely identifiable assets on the Ethereum blockchain and similar projects. The Ethereum unit ETH serves both as a exchangeable virtual currency unit analogous to BTC, but is also used as a utility token in order to pay verification nodes (called "miners") for the distributed execution of the general-purpose programs and smart contracts within the Ethereum ecosystem. This programmability and hundreds of alternative smart-contract-enabled blockchains have given rise to an entire ecosystem of complex lending, payment, and cryptocurrency exchange protocols, commonly referred to as decentralized finance – DeFi for short.

Due to the novel nature of the blockchain technology and the introduction of Bitcoin, the valuation of BTC in terms of established government-issued fiat currencies (like the US dollar) using traditional methods based on discounted cashflows or comparison to competitors is not applicable. Past efforts have been focused on aspects like scarcity, network effects and behavioral effects as well as certain technological details of the mining process. The valuation of smart contact enabled blockchain tokens adds the additional utility aspect of the currency unit when it serves as payment for the execution of smart contracts.

The goal of this work is to review past efforts and relevant considerations for the valuation of cryptocurrencies and show the relevance of DeFi metrics on smart-contract-enabled blockchains for this task.

The organization of this work is as follows: In chapter 2, the blockchain technology is reviewed, following the historical development of Bitcoin and Ethereum. The consensus mechanisms proof-of-work and proof-of-stake are discussed, as well as general-purpose decentralized application programmability via EVM smart contracts followed by the subsequent development of stablecoins, altcoins and decentralized finance. In chapter 3, several Bitcoin valuation methods are reviewed: the cost-of-production model, network growth and transaction volume based valuation techniques, the stock-to-flow model, as well as findings from the field of behavioral science. In chapter 4, several qualitative and conceptual considerations regarding the valuation of crypto assets are discussed, with a focus on legal uncertainties, occurrences of massive fraud and potentially cascading situations arising from the interdependence on stablecoins. The utility aspect of altcoins and applications of decentralized finance are also discussed. In chapter 5, a short-term Ethereum price prediction toy model based on machine learning techniques is introduced. Using different combinations of training data, it is shown that the next-day ETH price can be approximated based on historical BTC price and DeFi metric TVL data, but not on either input alone. This indicates that significant information regarding the smart contract altcoin price is encoded in the metrics of decentralized finance. Finally, in chapter 6, those findings are reviewed and summarized. The work concludes with several proposals for potential future research directions.

2 Technology of Cryptocurrencies

In this chapter the development of blockchain technology is summarized, starting from Bitcoin as the very first and still very active "blockchain 1.0" instance, Ethereum as the initiator of general-purpose smart contracts and the more recent development of altcoins, stablecoins and decentralized finance. Some of the technical aspects will be discussed in detail due to their importance for later valuation techniques, e.g., understanding proof-of-work vs. proof-of-stake consensus mechanisms and their associated energy usage.

2.1 Bitcoin

The blockchain technology underlying the Bitcoin network was first introduced in the original Nakamoto paper[6], which combines several established building blocks from the field of cryptography to solve the double spending problem and the related Byzantine generals problem[7].

A *cryptographic hash function* performs a deterministic computation, where an input of arbitrary size is being mapped to a fixed-length output, called the hash value, such that a minor change in the input (e.g., a single bit changed) leads to a significant change in the output that appears to be uncorrelated to the prior value. Furthermore, given a hash value it should be infeasible to generate an input that produces the hash value, i.e., the hash operation is supposed to be irreversible. In a similar sense, it should be infeasible to find two different inputs having the same hash value.[8]

Hash functions are typically used for integrity checks of large messages or as fingerprints to identify large data sets. In the context of the Bitcoin blockchain the infeasibility to reverse a hash value is used to construct a puzzle of adjustable computational complexity: a small additional amount of data, called a "nonce", is added after the main payload data. By requiring that the hash value of the entire data (payload and nonce) must end with a certain number of bits having the value zero, a

---

[6] See reference (Nakamoto 2008).

[7] The Byzantine generals problem describes the situation that multiple actors ("generals" in the historical example) need to act in a concentrated effort, but some of the actors are unreliable or act treacherously against the group. Byzantine fault tolerance (BFT) refers to the resiliency of a fault-tolerant system to such conditions, with Bitcoin being an important example.

[8] See section 5.1 of (Ferguson, Schneier and Kohno 2010).

brute-force search for the nonce is the only known approach to solve the puzzle due to the irreversibility of the hash function. Increasing the number of required zero bits in the output hash value increases the average number of required hash computations to find the nonce exponentially. This puzzle is used when data is added to the blockchain.[9]

Aside from hashes, the cryptographic technique of *digital signatures* is used to represent ownership. In asymmetric encryption schemes one party holds a pair of a private and a public key. The public key is being broadcasted out into the world, whereas the private key is kept in secret. Using the public key, another participant can encrypt a message, which can only be decrypted and read by the holder of the private key. In a similar fashion, but applied mathematically in reverse, a digital signature can be computed using the private key, whose validity can be checked and confirmed by everyone else using the broadcasted public key. Such private/public key pairs represent wallets, where the private key allows to access the funds by signing transactions, and a fingerprint of the public key serves as the wallet address when receiving money.[10]

The Bitcoin blockchain is consists of data blocks containing transaction information. Each transaction consists of a source and destination wallet address, the amount being transferred and is digitally signed using the private key of the sender's wallet. When a user submits a transaction to the Bitcoin network, it is broadcasted to numerous nodes and collected in a temporary buffer for unconfirmed transactions. Special nodes in the network, called *miners*, bundle several such transactions into a block, check all the signatures and account balances and include an additional special transaction that pays a certain amount to the miner's wallet, which is called the *mining fee*. The data block also includes the hash of the previous block in the blockchain, thus linking the block uniquely to earlier blocks. After the preparation of this candidate block completes, the search for the nonce starts. The Bitcoin network dynamically adjusts the complexity of the nonce puzzle, such that it takes on average about 10 minutes to solve the puzzle. Once a solution nonce for the candidate block has been found, the block and nonce are broadcasted into the network.

---

[9] See part 2 of (Franco 2014), chapter 10 of (Antonopoulos, Mastering Bitcoin 2017) and chapter 2 of (Lantz and Cawrey 2020).

[10] See section 12.7 of (Ferguson, Schneier and Kohno 2010).

Since the computation of a single hash for the proposed solution is rather fast, all participants in the network can almost instantaneously check if the puzzle has indeed been solved and can then accept or reject the new block after confirming the signatures and account balances once again. If everything checks out, the new block is appended to the blockchain, containing several thousand transactions. The added transactions are removed from the temporary buffers, the miners abandon their current search for a nonce, repackage new candidate blocks and the process begins again.

Originally considered in the early 1990s to combat spam email[11], this dynamically adjusted puzzle approach, called *proof-of-work*, provides a heartbeat to the entire blockchain network, which allows to synchronize the work. Naturally, with several independent miners working at the same time, the possibility arises that the puzzle for different candidate blocks will be solved at the same time. In such cases, the blockchain bifurcates, and the mining process continues for both chain heads with equal probability. However, since the longest chain is defined as the dominant primary chain, this situation quickly dissolves during the mining of the next blocks, as it is statistically highly unlikely, that a simultaneous solution for different blocks is found again. In the end, the mining work on one chain head wins the race and the new dominant chain is selected for subsequent mining work.

Ultimately, this system allows to establish a consensus between all participants and a transaction is considered to be settled once it is five blocks deep in the longest chain. Due to the chaining of blocks using the hash of the prior block, any chance of a bifurcation overtaking a longer portion of the chain must be able to solve the puzzle significantly faster than all the other participants. Since every miner wants to grab the mining fee for himself, which requires the mining fee transaction itself to be settled, the economic incentive for all miners is naturally to keep working on the chain, which is most likely to be the dominant one, i.e., the longest chain.

Miners provide the important service of verifying and confirming transactions, for which they are being rewarded by the *mining fee*. As part of the Bitcoin system, the mining fee for a candidate block is fixed and being halved every 210,000 blocks,

---

[11] See reference (Dwork and Naor 1993).

which according to the average 10-minute block intervals takes about four years. The difficulty of the mining puzzle is adjusted by the network every 2016 blocks, about every two weeks, to keep the average 10-minute interval. Those two weeks are referred to as a *mining difficulty epoch*.

At the time of writing (mid-2022) solving the nonce puzzle is being rewarded by 6.25 BTC plus the sum of all additional priority mining fees, that can be added to a transaction by the sender to speed up the confirmation, i.e., to be included in the next candidate block as soon as possible. Therefore, 6.25 BTC are newly minted every 10 minutes, with the next mining reward halving occurring in March 2024. Due to the halving schedule, the increase of the total BTC monetary unit supply asymptotically approaches 21 million BTC units, i.e., inflation goes to zero[12]. This also implies, that ultimately miners will be paid for primarily by priority fees once the mining fee alone becomes economically unsustainable. However, due to lost wallet keys, some BTC may be forever locked into wallets without any chance of ever being utilized again, thus the total usable amount of BTC monetary units may be significantly less than the mined supply.

The Bitcoin protocol also supports a limited form of smart contracts[13]. For example, a time-lock can be put in place, that prohibits the spending of BTC at the destination address until a certain date or time has been reached. Likewise, using multi-signature techniques, BTC can be locked in a transaction until signatures from multiple users have been provided.

In the historical context, the Bitcoin with its proof-of-work consensus mechanism and as the first widely accepted cryptocurrency is today viewed as a Blockchain 1.0 technology.

2.2 Ethereum

The restrictive scripting language that defines Bitcoin transactions and smart contracts was replaced by a general-purpose smart contract scripting language called

[12] This can be interpreted as a deflationary monetary system. See reference (Booth 2020) for an overview and the implications of a deflationary economic environment.

[13] This limitation has been gradually reduced over time, but a full general-purpose application environment analogous to the EVM is not envisioned for the Bitcoin ecosystem.

Solidity, which runs in the Ethereum Virtual Machine (EVM)[14]. Due to the Turing-completeness of this language, smart contracts cannot be automatically checked or decided, and execution becomes subject to the typical issues of general-purpose programming, e.g., the halting problem. Since the smart contracts must be executed and evaluated by the mining nodes as part of the process of adding a new block to the Ethereum blockchain, an economic incentive to spend the necessary computational resources must be installed, called a *gas fee*. Like the optional transaction priority fee in the Bitcoin ecosystem, the gas fee can be adjusted to deal with longer execution duration of smart contracts and to prioritize smart contract execution. Simple transactions like sending ETH from one wallet to another are specified in the same language as complex smart contracts and require gas to be executed as well. In essence, gas in the form of ETH powers the EVM and is therefore the lifeblood of the Ethereum ecosystem.

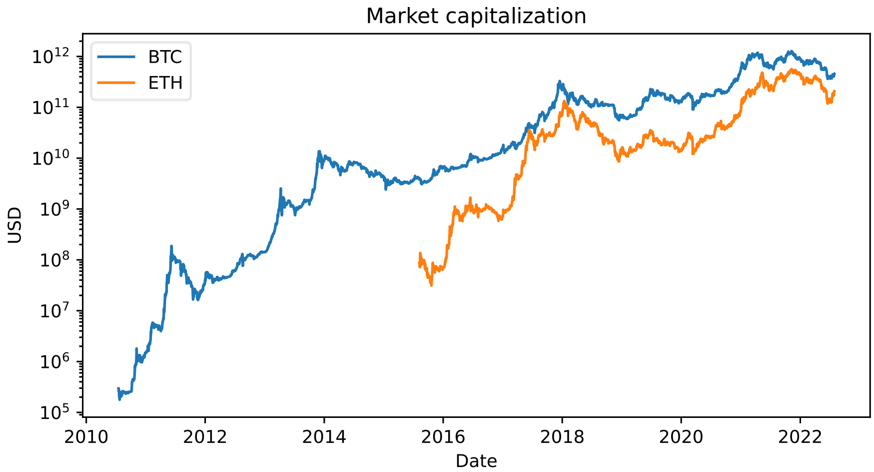

Figure 2: Market capitalization of Bitcoin and Ethereum
(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

Unlike in the Bitcoin ecosystem, the supply of ETH monetary units is not limited and is being expanded by 2 ETH per mined block every 12-15 seconds on average, such that since early 2019 around 13500 ETH have been added every day. However, since the so-called "London Upgrade" to the Ethereum network in late 2021, a part

[14] See references (Andres 2021), (Antonopoulos and Gavin, Mastering Ethereum 2018) or (Diedrich 2016) for a full introduction to the Ethereum blockchain and its technology. Reference (Russo 2020) provides a recollection of the development and first years of the Ethereum project.

of the gas fee depending on the current load of the network is being burned during the mining process, i.e., removed from circulation permanently. At the time of writing there are around 120 million ETH in circulation. Ultimately, due to the utility character of the monetary unit ETH as well as the different creation schedule and burning of ETH, the economy of the Ethereum ecosystem differs considerably from Bitcoin.

The Ethereum blockchain and Solidity smart contract language enable the implementation of non-fungible tokens (NFTs). Each NFT is a unique digital object and different from each other NFT. In contrast, all ETH monetary units are the same and simply add up to a total balance within each wallet, a property called fungibility. Due to the limited storage capacity in each block of the Ethereum blockchain, NFTs are typically implemented as a reference to a media file like an image, video or music file located on a different distributed storage medium, like the “inter-planetary file system” (IPFS). The programmability of digital assets using smart contracts allows to implement new forms of ownership and advanced transfer protocols for NFTs: for example, during the sale of an NFT, a royalty fee may be automatically paid to its original author.

Another application of the smart contract flexibility enabled by the Ethereum ecosystem are decentralized autonomous organizations (DAOs), where the classical hierarchy of a traditional organization as well as voting and spending powers of the leadership team are replaced by a smart contract, that implements a flat hierarchy and enforces a democratic system, where all members hold the same power. The membership is often implemented by usage of a governance token, i.e., a digital asset like an NFT that can be traded on decentralized exchanges. For example, the MakerDAO project maintains a cryptocurrency called Dai on the Ethereum blockchain and organizational membership is represented by the MKR governance token.

The general programmability of Ethereum with the subsequent development of distributed applications, complex smart contracts and NFTs is today referred to as a Blockchain 2.0 technology.

2.3 Further Altcoins

Aside from Bitcoin and Ethereum, several ten-thousand other blockchain-based cryptocurrencies are in existence. Many of those are obscure fungible tokens implemented on top of an existing blockchain network via smart contracts or represent governance tokens, but there are several large-scale projects that aim to improve blockchain technology itself, i.e., scalability issues like throughput, settlement, etc., and therefore run on their own blockchain networks.

The *Lightning network* implements a second layer payment protocol on top of the Bitcoin network[15]. Certain amounts of the BTC currency are locked into specific wallets, which makes those funds available to the Lightning network. Here transactions can be carried out in fractions of a second, which enables instant BTC payments without the need to wait for regular Bitcoin settlement. After certain intervals, only the net transactions facilitated within the Lightning network are being written back to the main Bitcoin blockchain, which gives rise to the notion of "on-chain" and "off-chain" transactions.

Many of the altcoins supporting general-purpose programmability utilize the Ethereum Virtual Machine and therefore allow to reuse smart contracts on different blockchains, though there are notable exceptions like Cardano, for example. One application is to implement *2nd-layer networks* on top of an existing blockchain. For example, the Polygon blockchain with the MATIC fungible token is one of several competing EVM-compatible scalability solutions, that aims to deal with the limited throughput of the Ethereum network and the considerable gas fees during times of network congestion.

Alternative blockchains also attempt to reduce the issue of energy wastage during the mining process of the proof-of-work consensus mechanism. While the hash-based cryptographic puzzle and its dynamically adjustable complexity solve the double spending problem, an increase of the network capacity and computational resource ultimately result in more energy being spend on the mining process. This energy is effectively wasted, as the solutions to the artificial cryptographic puzzle are useless except in the context of blockchain confirmations. Thus, technological

[15] See reference (Antonopoulos, Osuntokun and Pickhardt, Mastering the Lightning Network 2021) for a full introduction to the Lightning network and its technology.

progress and higher efficiency due to improvements in semiconductor technologies are effectively negated by increasing the puzzle difficulty to keep the block verification times steady. While originally negligible, the economic incentives in mining led to the situation, that the Bitcoin network alone in mid-2022 consumes about 130 TWh of energy annually, more than entire countries. A single transaction in the Bitcoin network consumes almost 1500 kWh of power, which is incredibly inefficient compared to traditional centralized payment solutions[16]. Aside from the energy usage, the proof-of-work approach also incentivizes centralization due to economics of scale, contrary to the original goal of a decentralized payment system: since mining rewards are only paid out to the sole ultimately settled winner in the race to solve the cryptographic puzzle, the efficiency gains of a large-scale mining operation and pooling of resources leads to a reduction in the number of independent miners.

Instead of the proof-of-work consensus mechanism, alternatives like *proof-of-stake* are being used by more and more blockchains. Instead of solving a cryptographic puzzle, a validator node stakes a number of tokens and is randomly chosen to validate a new block with the probability of being chosen proportionally to the amount staked. Being selected allows the validator to collect the transaction fees associated with the verification process of the block. If a validator approves a fraudulent transaction, the staked amount will be lost, i.e., the economic incentive to be chosen more frequently by staking larger amounts of tokens also incentivizes to act in a trustworthy manner. The main benefit of this approach is that it avoids the energy wasted during the solving of the cryptographic puzzle and leads to more decentralization, as the economy-of-scale benefits of larger mining operations are avoided. At the time of writing, the Ethereum blockchain currently undergoes a lengthy process of changing from the proof-of-work consensus mechanism to a more scalable and energy-efficient proof-of-stake approach, which is often referred to as Ethereum 2.0.

With the increasing number of alternative blockchains, the requirement to transfer assets between different blockchains arises. For example, while numerous blockchains with general programmability support NFTs, an NFT is only unique with

[16] See reference (Bitcoin Energy Consumption Index 2022).

respect to its respective blockchain. Several projects like Cosmos or Polkadot try to address this issue by developing an inter-blockchain infrastructure, that supports the exchange of assets between different blockchains.

Improvements in the scalability of blockchain settlement, throughput and energy-usage combined with inter-blockchain technologies is today referred to as Blockchain 3.0 technology. A noteworthy upgrade is the Ethereum 2.0 transition to the proof-of-stake consensus mechanism and various other improvements, which marks the only large-scale project undergoing a Blockchain 2.0 to 3.0 technology upgrade.

2.4 Stablecoins and Decentralized Finance

Aside from mining, the commonly used method to obtain monetary units of any cryptocurrency is to trade it on an exchange. Well-known cryptocurrency exchanges like Coinbase or Kraken are registered companies that allow to exchange government-issued fiat currencies for various cryptocurrencies. Those are centralized entities, that represent a bridge between the traditional fiat monetary system and the blockchain-based cryptocurrencies.

With the rise of the general-purpose programmability introduced by the Ethereum blockchain, the idea of a decentralized exchange (DEX) was introduced, where the actual exchange of one cryptocurrency to another is entirely facilitated within transparent and user-inspectable smart contracts on the respective blockchains. Such decentralized exchanges are part of an entire ecosystem of borrowing and lending providers, which are being enabled by users voluntarily staking their cryptocurrencies into designated pools to earn interest, similar to a traditional bank's lending of customer savings while paying an interest. This gives rise to the notion of *decentralized finance* (DeFi), where most activities of the banking industry are replicated in terms of smart contracts and the middleman is replaced by code[17]. For example, the exchange of one cryptocurrency to another on a DEX is facilitated by automated market makers (AMMs), where the user effectively trades with the smart contract.

[17] See reference (Harvey, Ramachandran and Santoro 2021) for a comprehensive overview of the DeFi ecosystem and current applications.

Trading to exploit short-term opportunities in the highly volatile cryptocurrency markets requires an efficient access to a stable currency, like the US dollar, for example. However, the bridge to classical fiat currencies is typically rather slow (on the order of hours or days) and only possible using centralized exchanges. This gave rise to the idea of a *stablecoin*, a cryptocurrency designed to keep a close price to a reference currency like the US dollar, such that other cryptocurrencies can be traded in stablecoins on DEXes without being exchanged into actual fiat currencies. In 2014 the very first stablecoins BitUSD and NuBits were introduced, followed by Tether USDT a year later from the Hong Kong-registered cryptocurrency exchange Bitfinex. All three of those projects are still active today, even though both BitUSD and NuBits have lost their original peg to the US dollar—the first one trading at 0,82 USD and latter one trading at 0,00036 USD in mid-2022.

The majority of (high frequency) trading in cryptocurrencies today relies on the usage of stablecoins. The peg can be realized by fiat-backing, where a third-party regulated company issues a promise to pay the guaranteed amount on demand, e.g., USD Tether, or by cryptocurrency backing, where a basket of other supposedly valuable cryptocurrency is being kept as collateral that can be sold when a certain amount is being redeemed, e.g., the DAI stablecoin issued by the MakerDAO project. Furthermore, algorithmic stablecoins utilize purely automated trading to control a stablecoins money supply and its value, where no collateral is being kept in reserve. Recently, collapses like the Terra/Luna UST stablecoin in May 2022, which wiped out around $45 billion of market capitalization within a couple of days, have put the feasibility of such algorithmic stablecoins without any actual backing to independent assets into question.

2.5 Futures and other crypto derivatives

Futures are standardized legal contracts, which handle the purchase or sale of a certain asset, like copper, wheat or stocks, at a specific time in the future and for a certain predetermined price between unknown parties at the time of writing. On futures exchanges, where such contracts are being traded, buyers and sellers are matched, and transactions are settled once the expiry date is reached. With changes in value of the underlying asset, the value of the futures contract changes accordingly.

The first Bitcoin futures were listed in December 2017 on the Chicago Mercantile Exchange (CME) and Chicago Board Options Exchange (CBOE), which are settled in US dollar cash upon expiry. On such regulated exchanges, the margin requirement for investors in BTC futures is around 50%, providing a leverage factor of 2. On unregulated exchanges like Binance, a leverage factor of up to 125 was allowed until July 2021, which considering the highly volatile nature of cryptocurrencies in general led to excessive risk taking by traders. In January 2020 CME introduced cryptocurrency options build on top of the cryptocurrency futures, a construction necessary for regulatory reasons. In October 2021 the first Bitcoin ETF build on futures was approved for the US market, since regulatory reasons still prohibit a regulated fund from direct investment into cryptocurrencies.

The usage cash-settled futures and other derivatives on regulated exchanges offers the possibility of cryptocurrency exposure for investors with established financial instruments and without using the blockchain technology, which is an important aspect especially for institutional investors that want to diversify without additional technological burdens.

3 Quantitative valuation models for Bitcoin

The first transaction that put a real-world value context to the novel digital currency BTC was carried out on May 22nd, 2010, sixteen months after the mining of the very first "Genesis" block of the Bitcoin blockchain. The delivery of two pizzas was being paid for with 10000 BTC by the crypto enthusiast Laszlo Hanyecz, roughly worth 40 US dollars. This occurred two months prior to the launch of the now defunct first cryptocurrency exchange Mt. Gox, that was based in Japan and dominated the market from 2010 to 2014, handling 70% of all cryptocurrency trades at its peak before abruptly shutting down after the loss of massive amounts of customer funds due to hacks and thefts became public. In the very first trade on Mt. Gox 20 BTC was exchanged for 0,99 USD, putting the value of BTC at around 0,05 USD/BTC on July 17, 2010. The collapse of Mt. Cox caused a crash in the cryptocurrency market with the Bitcoin price going down from more than 1000 USD/BTC to 200 USD/BTC in January 2014. Thus, in less than four years, the price being paid for a single monetary unit of Bitcoin did rise by five orders of magnitude from 0,004 USD/BTC at the time of the historic pizza payment. At the time of writing in mid-July 2022, the price for one BTC has come down to 20-22000 USD, about a 70% drawdown from its most recent all-time high[18] of 69044,77 USD on November 20, 2021.

Following the meteoric rise in value, the question arises how the fair price of this highly volatile asset class can be determined and ideally be forecasted. For traditional assets widely accepted models for estimating future prices exist based on fundamental data. The discounted cashflow model is one reference model used to value companies. For bonds, a systematic analysis of debt-to-equity levels and credit ratings give reasonable risk/reward estimates for an investment. And to value real estate, a comparison to the neighborhood or similar situated properties can be considered.

With respect to cryptocurrencies, especially the case of Bitcoin considered here, all those methods are not applicable. The notions of cashflow, debt, equity, credit rating, or a comparison to similar assets are ill-defined for the BTC monetary unit.

[18] Data retrieved from CoinGecko.com.

While short-term based trading tools like technical analysis and momentum strategies have been successfully applied since the early days of Mt. Gox, a fundamental approach appears to be elusive—there is no consensus how to value Bitcoin or other cryptocurrencies appropriately. In the following several methods for valuation of the BTC monetary unit are reviewed.[19]

3.1 Cost-based estimates

One of the first systematic approaches to cryptocurrency valuation is Adam Hayes's cost of production method[20], which attempts to value a BTC monetary unit based on the cost of production of this unit, i.e., the variable costs of the electricity spent during mining and the fixed costs of the mining infrastructure. Assuming an economically rational acting miner, the mining itself will only take place if the cost of electricity, the current level of mining difficulty and the market price of a BTC unit make this a profitable endeavor. Using a current and widely used Bitmain Antminer S19 Pro ASIC system, the current price of around 20-22,000 USD/BTC in mid-July 2022 is right at the profitability threshold.[21]

The cost of production model is only applicable to cryptocurrencies that utilize the proof-of-work consensus mechanism. More and more blockchain projects are built on the more energy-efficient and scalable proof-of-stake consensus, including established projects progressing towards proof-of-stake like Ethereum. Furthermore, additional transaction fees and their impact on miner profitability are not considered. The model also assumes that all miners are in a competitive situation and ignore the effects of economies-of-scale centralization, which led to the oligopolistic situation that a small number of large-scale mining operations dominate the total mining power.

Hayes's method belongs to the absolute pricing models, as an explicit BTC/USD price estimate can be computed. Today it is typically considered to be somewhat of

[19] Partially based on reference (Ige 2018).
[20] See reference (Hayes 2015).
[21] The Antminer S19 Pro system computes 110 TeraHashes per second and draws 3250W of electricity. Profitability is calculated at an assumed global average energy price of 0,12 USD/kWh. See https://www.asicminervalue.com/miners/bitmain/antminer-s19-pro-110th for details.

a lower bound for value estimates due to the aforementioned ignored aspects that further raise the value.

3.2 Network valuation models

Another approach to value Bitcoin derives from its design goal as a peer-to-peer payment system, which becomes more useful with a growing number of participants. The network size and growth become the base metrics for new valuation models. This follows roughly the valuation growth of Internet-based companies, which has been strongly linked to a growing user basis especially in the early years due to both internal scaling effects and increasing utility for interacting users. This effect has been referred to as *Metcalfe's law* in computer science, named after an early network engineering pioneer of the internet revolution from the 1970s, who proposed that the value of a network increases quadratically with a growing number of network nodes, i.e., the value of a network in relative terms is proportional to the number of possible connections between the participants.

In a 2006 paper titled "Metcalfe's Law is Wrong"[22], the authors argue that the quadratic increase in network value may be too large and rather follows an $n \log n$ scaling behavior, which has been referred to as *Odlyzko's law* subsequently. Despite being put forward by Metcalfe in the early 1980s, the discrepancy remained largely unsettled due to a lack of data. A team of Dutch researchers analyzed European internet usage data[23] in 2013 and found that Metcalfe's n^2-law holds true for smaller network sizes n and changes to Odlyzko's $n \log n$-law for large n, in essence validating both scaling models for different network sizes. Metcalfe himself provided further evidence for his law by analyzing Facebook data[24], which was independently strengthened by a subsequent analysis[25] based on data from internet giants Tencent and Facebook in 2015, ultimately leading to strong empirical evidence for the n^2-law for network growth value in relevant domain sizes.

[22] See reference (Briscoe, Odlyzko and Tilly 2006).
[23] See reference (Madureira, et al. 2013).
[24] See reference (Metcalfe 2013).
[25] See reference (Zhang, Liu and Xu 2015).

Following some earlier work by Ken Alabi[26], both Peterson and Van Vliet in 2018 applied Metcalfe's law to Bitcoin valuation[27] and showed a correlation of more than 80% between network growth and BTC price over medium to long-term time horizons. Two years later, the network size to BTC price correlation was further investigated and found to explain more than 90% of the Bitcoin market capitalization[28].

The network size in those works is approximated by the number of distinct Bitcoin addresses[29], which can be automatically determined by scanning the Bitcoin blockchain. However, this is a rather rough measure for the number of distinct participants in the BTC payment network, as users can have multiple wallets and account balances are often spread over multiple cryptographically linked addresses to obfuscate ownership and enhance payment pseudo-anonymity. On the other hand, exchanges often only use a rather small number of addresses. Based on the identified correlations to Metcalfe's law, both overcounting and undercounting appear to cancel each other.

In general, such pricing forecasts are referred to as "Network Value to Metcalfe's Law" (NVML) or "Network Value to Odlyzko's Law" (NVOL) models. Both are relative pricing models, as a scaling behavior aspect is utilized to extrapolate current BTC/USD prices into the future based on changes in the network size.

3.3 Network Value to Transaction

To create a metric akin to the price-to-earnings (P/E) ratio for classical stocks and companies, Woo introduced[30] the "Network Value to Transactions" (NVT) ratio in 2017, which is defined as the total market capitalization divided by the total daily transfer volume. Like the P/E ratio, a high NVT ratio indicates a potential overpricing of the cryptocurrency, whereas a low NVT ratio suggest an entry point for purchases. The NVT ratio attempts to connect the store of value character of Bitcoin

[26] See reference (Alabi 2017).
[27] See references (Peterson 2018) and (Van Vliet 2018).
[28] See reference (Cipolaro and Stevens 2020).
[29] Bitcoin addresses consist of a 160-bit hash of the wallet public key, i.e., there are potentially $2^{160} \approx 1.46 \cdot 10^{48}$ wallet addresses. The space of Bitcoin addresses is therefore sparsely populated and could be expanded in the future through upgrades.
[30] See reference (Woo 2017).

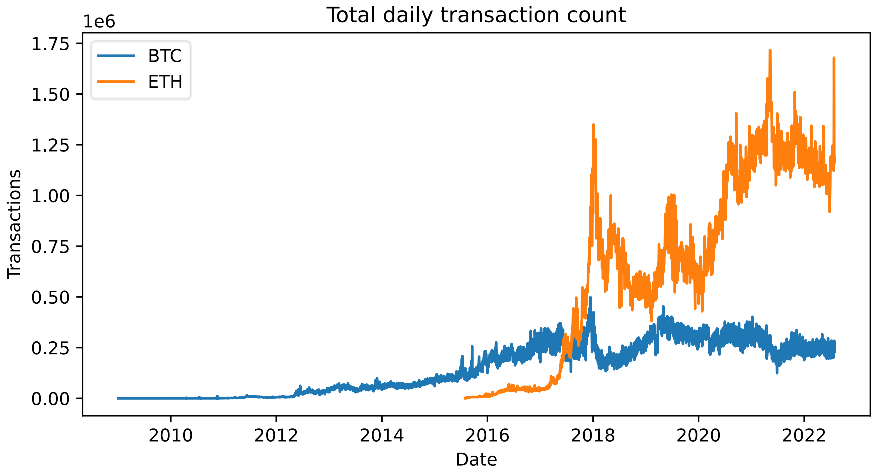

Figure 3: Total daily Tx count on the Bitcoin and Ethereum blockchain
(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

to its utility aspect as a decentralized payment system, with the latter being measured by the transactional throughput. While the monetary transactional value in unrestricted, the Bitcoin and Ethereum blockchain technology is currently limited to a relatively modest 6-7 and 20-25 transactions per second, respectively, which can be seen in the relative flatness of the daily transactions seen in Figure 3.

The NVT ratio is referred to as an on-chain analysis tool, as the data for the transfer volumes is automatically extracted from the blockchain. With the implementation of the Lightning network to handle instantaneous low-cost transactions, parts of the transactional volume become invisible from the Bitcoin blockchain, such that the naïve NVT on-chain analysis becomes more and more distorted. Like the network value models, the NVT ratio is a relative pricing instrument by considering historical averages and deriving an associated BTC price based on current transactional volume.

3.4 Stock-to-flow model

One of the most discussed forecasting models for the Bitcoin price is the "stock-to-flow" (SF or S2F) model[31], which uses a measure for scarcity as its fundamental valuation principle. It was put forward in 2019 by a Dutch institutional investor with the name PlanB and is motivated by a characterization of the intrinsic value of

[31] See reference (PlanB 2019).

scarcity by Nick Szabo[32], an American cryptographer and computer scientists, preceding the release of the original Nakamoto paper. Szabo writes in his blog:

> "What do antiques, time, and gold have in common? They are costly, due either to their original cost or the improbability of their history, and it is difficult to spoof this costliness. [..] There are some problems involved with implementing *unforgeable costliness* on a computer. If such problems can be overcome, we can achieve bit gold."

> "Precious metals and collectibles have an *unforgeable scarcity* due to the costliness of their creation. This once provided money the value of which was largely independent of any trusted third party. [..][but] you can't pay online with metal. Thus, it would be very nice if there were a protocol whereby unforgeably costly bits could be created online with minimal dependence on trusted third parties, and then securely stored, transferred, and assayed with similar minimal trust. Bit gold."

Following Szabo, the unforgeable scarcity due to the limited amount of BTC units and unforgeable costliness due to difficulty required for mining new BTC units link Bitcoin to scarce elements like gold and silver. Within this precious metal universe a measure called the stock-to-flow ratio has previously been considered for goods with significant stock quantities in order to relate current relative mining production to historical averages and to provide a measure for the scarcity[33]. It measures the years required at current production levels to rebuild the currently existing supply. In the case of gold and silver the historic mean stock-to-flow ratios from 1900 to 2013 are 67 and 74, respectively, which have dropped to 2019 levels of 62 and 22 due to increased mining activity and rising prices in recent years. Furthermore, a non-perishable good with a high stock-to-flow ratio has been associated with its suitability as a monetary unit and store of value.

---

[32] See references (Szabo, Antiques, time, gold, and bit gold 2008) and (Szabo, Bit gold 2008).

[33] See reference (In Gold We Trust Report 2012).

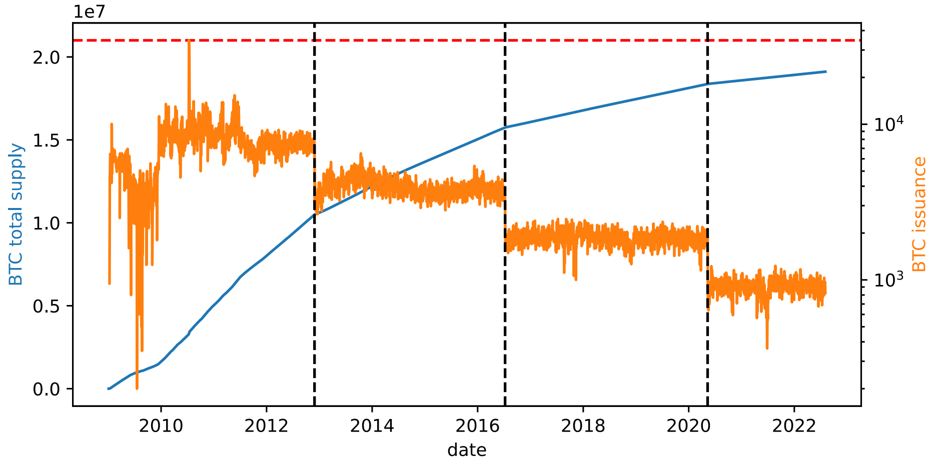

Figure 4: Bitcoin BTC issuance and total supply
(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

Applied to the cryptocurrency space by Ammous[34] and further developed by PlanB, the stock-to-flow ratio has been investigated for Bitcoin. Due to the halving of mining rewards in the proof-of-work consensus, the quantity of newly issued coins decreases over time, which in turn increases the scarcity of the BTC monetary unit, compare Figure 4. Such Bitcoin mining reward halving events have historically caused a big surge in the BTC/USD price, with the next one taking place in March 2024.

Due to the fixed schedule of the Bitcoin halvings, its current stock of around 19.1 million mined BTC units and an annual mining production of 328500 new BTC units leads to a current stock-to-flow ratio of around 58, i.e., mining at the current rate would require 58 years of mining to reach the current total BTC unit supply. PlanB then applied the stock-to-flow ratio in a regression to derive an absolute pricing model:

$$\text{S2F:}\quad \text{USD/BTC} = 0.4 \cdot \text{SF}^3$$

This model follows a simple cubic scaling law of the stock-to-flow ratio. The basic methodology has been confirmed by independent sources[35] with minor variations in the coefficients.

---

[34] See chapter 3 of (Ammous 2018).

[35] See https://medium.datadriveninvestor.com/stock-to-flow-summarized-the-code-a-brief-outline-and-its-limitations-42a524163dde for an explicit derivation.

Due to its popularity and structural difference from other models, there has also been a variety of critiques to the S2F model. It has been noted that the S2F value of gold is varying significantly over time and appears to be uncorrelated to the price of gold, i.e., a modelling of the gold price based on the same approach appears to be rather unsuccessful. Likewise, the artificial scarcity of Bitcoin has been put into question with the release of thousands of alternative blockchain-based coins sharing similar designs. Whereas physical materials like gold and silver have fundamentally unique properties that cannot be replicated, the digital scarcity of BTC is only enforced within the Bitcoin blockchain itself. Its properties, on the other hand, have been duplicated numerous times in various obscure altcoin projects. The stock-to-flow model fails to recognize lost or inaccessible funds due to lost keys or other causes[36], which reduces the relevant stock. Significant quantities of BTC mined during the first years of the Bitcoin blockchain, currently worth tens of billions of US dollars, have never been moved and may be lost forever from further usage.

Furthermore, several statistical and technical problems of the model[37] have been pointed out, and multiple adjustments to the model parameters for an improved fit after the initial publication have been criticized. An in-depth statistical analysis of the S2F model by Burger[38] initially rejected the usage of a linear regression for modelling the BTC price based on the S2F ratio due to a lack of key statistical assumptions. After considering the special case of cointegration between the logarithm of the stock-to-flow ratio of Bitcoin and the logarithm of its price, which would indeed allow for the regression approach as published, it was ultimately found that both time series of the BTC price and S2F ratio do not share an equal order of integration. Therefore, the statistical basis for the stock-to-flow model appears to be invalidated, despite its widespread popularity. Yet, to this date, the stock-to-flow model remains the only absolute Bitcoin pricing model that replicated

---

[36] Sending BTC to a valid yet unoccupied wallet address makes the funds inaccessible and they are lost for any future transaction. This method is sometimes used to voluntarily burn funds.

[37] See the conference talk (Kripfganz 2020).

[38] See reference (Burger, Challenging PlanB: a review of modelling Bitcoins value with scarcity 2019) for the initial criticism, (Burger, Reviewing “Modelling Bitcoin’s Value with Scarcity” —Part II: The hunt for cointegration 2019) for the first cointegration analysis and (Burger, Reviewing “Modelling Bitcoin’s Value with Scarcity” — Part III: The Fall Of Cointegration 2020) for the final criticism invalidating the S2F model.

the order-of-magnitude price increments with an empirically fixed delay from the mining reward halvings.

3.5 Behavioral approaches

Models to capture human behavior with regard to the buying and selling of assets have been considered in traditional finance for some time. In the cryptocurrency space, especially regarding Bitcoin, the research is focused primarily on bubble formation and bursting processes. The volatile nature of the crypto market in general has gone through relatively short-lived boom and bust phases, with skyrocketing price increases followed by 80-90% retractions in just a few months. Those so-called "crypto bubbles" and "crypto winters" are often stirred by social media and influencers. For example, a social media phenomenon called "hodling" refers to the behavior of never selling crypto assets and "holding on forever", following the belief, that the value of such assets will rise forever. The prevalence of such behavioral anomalies, especially in the case of retail investors, has been a characteristic of the cryptocurrency space since the very beginning.

Furthermore, significant differences have been found in the behavioral structure between Bitcoin and Ethereum investors[39], which react differently in situations of local price fluctuations and large systemic events. Bitcoin users tend to take a short-term view of the market in case of local situations but show a more optimistic outlook in case of larger events, whereas Ethereum users have the opposite reactions. This may be because Bitcoin as the oldest and most capitalized blockchain ecosystem provides a sort of stability center within the entire crypto space. It could also signal the structural differences between the "digital gold" Bitcoin and the development platform character of Ethereum with its general-purpose programmability and secondary usage as a utility.

[39] See reference (Aspembitova, Feng and Chew 2021).

4 Non-quantitative valuation aspects of crypto assets

Attempting to develop a valuation model for Bitcoin already proves to be a challenging task, as seen in the previous chapter, even though it is both the oldest and technologically least sophisticated instance of a blockchain network. Implementing general-purpose programmability and the subsequent emergence of decentralized finance are properties that turn the associated cryptocurrency into a utility aside from being a storage of value.

Further adding to the complexity of the valuation problem, the topic of illegal and fraudulent activities surrounding the crypto space cannot be ignored. This has already heavily influenced pricing levels in the past and there are several potentially devastating downfalls still waiting to happen, especially in the area of stablecoins. In this chapter further considerations relevant for the development of a more suitable valuation model are discussed.

4.1 Censorship-resistance

Blockchain technology provides strong censorship resistance from governmental intervention due to its fundamentally distributed nature. There are no central elements that can be easily taken over to control the network. Any major direction changes or data modifications on the blockchain itself must be convened by a majority decision[40] of the network nodes. Furthermore, the technological nature of wallet addresses in the form of key hashes provides a certain level of pseudo-anonymity for the network users. An address on their own does not reveal any information

[40] Majority attacks on any large-scale blockchain network are highly unlikely, as this would undermine the value of the very asset that one would attempt to steal. Furthermore, finding consensus for such a major change has historically been very difficult: The very first large-scale DAO build on the Ethereum blockchain (called "The DAO", a crowdfunding project supporting the further development of Ethereum) was hacked and around 60 million USD worth of ETH was stolen due to a coding error in the underlying smart contract. This led to a hardfork of the Ethereum blockchain where the fraudulent (yet technically valid) transactions were reversed. To this day the original Ethereum blockchain continues under the name "Ethereum Classic", but most miners have put their efforts behind the modified chain. Another example for such a majority decision was a proposal to increase the block size of the Bitcoin blockchain to improve throughput and scalability, leading to significant disagreements and discussions between core programmers and miners, see the book (Bier 2021) for a detailed account.

about the owner, but transactions facilitated by users with associated known addresses can be easily traced on a non-privacy-focused blockchain like Bitcoin or Ethereum.

Real-world links between users and wallets are typically established from the usage of regulated exchanges, which either require government-issued identity information or can indirectly obtain such data from fiat currency payment methods, e.g., credit cards or wire transfers. Once de-anonymized, every past or subsequent transaction involving the wallet address can be traced back to the user. The traceability issue of Bitcoin or Ethereum transactions has been technologically addressed by other blockchain projects centered around anonymity, e.g., privacy coins like Monero, Dash, zCash, etc., which use advanced cryptographical means to make transactions difficult to trace[41].

During the Canadian trucker protests against pandemic restrictions in early-2022 the government froze numerous bank accounts, leaving the protesters in need for essential items. Following the original idea of a censorship-free payment system, the truckers received numerous donations in various cryptocurrencies from supporters all over the world after a GoFundMe.com-based crowdfunding campaign was suspended and frozen as well. Several donation wallets on different blockchains have been placed on watch lists, such that a conversion of the donations back into fiat currency on a regulated exchange becomes virtually impossible without raising red flags. In addition, the wallet addresses of supporters are also transparent on most chains, putting them in potentially problematic situations via tracebacks as well. In a similar attempt to circumvent international monetary sanctions against Russia after the start of the Russia-Ukraine war in 2022, Russian hints to accept cryptocurrency payments for fossil fuel delivery put companies in an equally awkward position.

Regarding the valuation aspect of cryptocurrencies, one of the most proclaimed benefits therefore turns out to be at least partially wrong. While pseudo-anonymity and censorship resistance are indeed provided by the technological base layer, both

[41] While privacy coins proclaim the statement that transactions are effectively untraceable, the past has shown that it is possible to breach the privacy layer, see (Möser, et al. 2018).

are ineffective in current real-world scenarios—even in the case of so-called privacy coins. One of the original incentives that led to the publication of the Nakamoto whitepaper was to provide the world with a censorship-free monetary alternative, which is not subject to governmental or central bank intervention. Considering the anonymity intrinsic to the physical exchange of gold or cash payments, most blockchain-based payment networks currently lack this property. Both aspects are difficult to value at a quantitative level but should be recognized when considering "the case for crypto" at a fundamental level.

4.2 Cryptocurrency usage for illegal activities

Since the early days of Bitcoin, a significant portion of cryptocurrency transactions has revolved around illegal activities. For example, Bitcoin was initially the only method of payment on the dark net marketplace "Silk Road", where illegal drugs of all kinds, fake drivers licenses and other questionable items and services were traded from February 2011 to October 2013, until its founder was arrested by the FBI. To this day, ransomware attacks on company networks and cybercrime in general are typically required to be paid for in cryptocurrencies. The privacy coin Monero has become the blackmailer's favorite choice of money.

To hide transactions involving illegal activities, so-called cryptocurrency tumblers or mixers have been developed early on[42] for money laundering purposes, where cryptocurrency balances of multiple users are accumulated in the same wallet address and are then redistributed (minus an appropriate fee for the service) in random cuts of the original amount to numerous other destination addresses controlled by the original users. Considering the traceability discussed earlier, the idea is to mix illegal and legal transactions. Other services combine transaction mixers, unregulated exchanges and intermediate dummy transactions using privacy coins, which makes tracing payments rather difficult. Together with the censorship resistance, cryptocurrencies therefore provide a rather convenient payment ecosystem for criminal activities, complemented with an entire ecosystem for money laundering and traceability obfuscation[43] if done properly.

---

[42] For example, the crypto tumbler "Bitcoin Fog" was at least ten years in operation, when its founder Roman Sterlingov was arrested in April 2021 by the IRS.

[43] See reference (Europol 2021).

While it is very hard to quantify the number of transactions or total volume involving illegal activities in absolute terms, many governmental cybercrime organizations agree that it makes up a rather significant portion and that crypto has become the primary means of payment for illegal activities. Considering the issue of valuation, this taint presents a significant entry barrier for investors, especially from the institutional side, which simply keep a distance from the crypto space due to its association with criminal activity.

4.3 Stablecoin instabilities

Originally introduced to negate the highly volatile nature of the cryptocurrency market, stablecoins have become major cornerstones at exchanges, in decentralized finance and the entire crypto space[44]: in mid-2022 the three largest stablecoins Tether (UDST), USD Coin (USDC) and Binance USD (BUSD) are ranked 3rd, 4th and 6th in term of total market capitalization. This represents a total of around 140 billion USD in value or about 14% of the total cryptocurrency market, whereas just 18 months earlier in December 2020 stablecoins contributed only 2% to the total crypto market capitalization. With 45 billion USD in daily trading volume Tether ranks first in terms of activity, which is more than twice the volume of Bitcoin traded during the same 24-hour time window and seven times the volume of the second largest stablecoin USD Coin[45].

Due to this level of usage and importance in trading and decentralized finance, any sign of instability or unpegging can easily lead to massive volatility swings in the crypto space and trigger mass liquidations, followed by a massive market crash[46]. Furthermore, whereas the market capitalization of a cryptocurrency is somewhat fictitious[47], fiat-backed stablecoins represent actual money, i.e., the current market

---

[44] See the table on top-10 cryptocurrency market information in the appendix.

[45] Data retrieved from CoinMarketCap.com.

[46] A significant depegging of a stablecoin is quickly being exploited by high-frequency trading bots looking for arbitrage opportunities. Other bots are designed around the expectation of a stable peg, such that those automated actors together can quickly destabilize the market. Compare section 4.4 for related details.

[47] The current market capitalization is derived from the current market price, just like for shares of publicly traded companies. Using liquidity pools and borrowing protocols, a cryptocurrency with minimal fiat backing can be blown up several orders of magnitude. Furthermore, the actively traded cryptocurrency portion is rather small for most coins other than Bitcoin or Ethereum, which is one of the reasons for the highly volatile behavior.

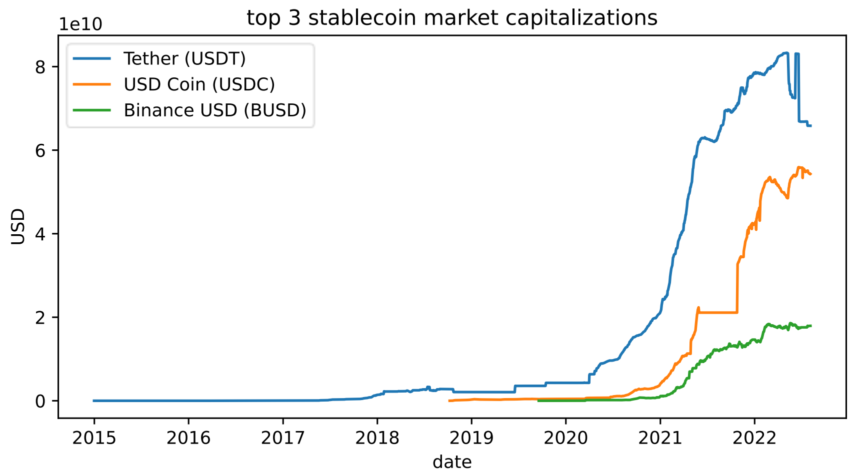

Figure 5: Market capitalization of Tether, USD Coin and Binance USD
(Created using data retrieved from CoinCodex, see https://coincodex.com.)

capitalization of the top-3 stablecoins is treated to be readily available for fiat currency redemptions and withdrawals. Following the collapse of the algorithmic stablecoin Terra/Luna in May 2022 with the loss of 45 billion USD in value, fiat-backed stablecoins have also come under more scrutiny by government regulators due to the uncertainties in their backings[48].

At the time of writing, the second largest stablecoin USD Coin (USDC) appears to be fully backed by reserve assets. It was released in 2018 by Circle and is managed by a consortium including members of the regulated cryptocurrency exchange Coinbase and the Bitcoin mining hardware manufacturer Bitmain. The USDC reserves are regularly attested by the regulated accounting firm Grant Thornton, LLP[49], however, those are not audits like for financial assets of regulated companies. When tracking the balance statements of the two primarily associated banks, Silvergate Bank (SI) and Signature Bank of NY (SBNY), both of which are publicly traded and have 24/7 transaction support to deal with the technological requirements of the never-sleeping crypto space, one can indeed recognize an appropriate rise in the respective balance sheets. Still, there is a significant bank run risk due to the size of around 50 billion USD in backing funds, which in theory must be readily available for immediate redemption.

---

[48] See reference (Ranger 2022).
[49] See https://www.centre.io/usdc-transparency for the monthly reports on the Centre Consortium's website.

Tether, on the other hand, appears to be a ticking time bomb and has been for years[50]. It has never been properly audited and only presented attestations from an unknown accounting firm with apparently only a single professional accountant on staff. In its early years, numerous entanglements with the unregulated and unbanked cryptocurrency exchange BitFinex, which enabled Tether trading in early 2015, did raise suspicions. The Paradise Paper leaks revealed that Bitfinex officials set up the holding company for Tether on the British Virgin Islands, despite claiming BitFinex and Tether to be two separate entities[51]. During the 2017 crypto boom, when Bitcoin briefly reached almost 20,000 USD, the amount of issued USDT grew from 10 million to 2.8 billion and subsequent research strongly suggests that half of the price increase was due to a manipulation scheme involving Tether[52]. Furthermore, Tether has declared suspiciously unchanging excess assets of around 160 million USD during the last two years (mid-2020 to mid-2022), despite the total amount of USDT growing from 10 billion to 84 billion at its peak and reducing to 66 billion at the time of writing. The only bank doing business directly with Tether is Deltec Bank from the Bahamas, whose balance sheets do not reflect the growth of more than 70 billion US dollars in equivalent assets. As revealed by the „Tether Papers“[53], almost half of this Tether growth was due to minting requested by the two high-frequency cryptocurrency trading companies Alameda Research and Cumberland Global, the former of which is deeply entangled with the cryptocurrency exchange FTX and both are currently facing market manipulation allegations[54]. Investigations suggest that Tether and Bitfinex are using each other’s funds to cover losses and redemptions like in a single entity, i.e., Bitfinex losses in the past have been covered by Tether funds and Tether redemptions seem to be serviced by Bitfinex customer fund withdrawals.

During the collapse of the 45 billion USD stablecoin Terra/Luna in May 2022 there were first signs of a depegging of Tether, when it briefly traded at 0.9485 US dollars, but quickly bounced back. Tether would collapse as soon as it is unable to service a redemption due to a lack of actual funds. Considering the daily trading

[50] See reference (McKenzie 2022).
[51] See reference (Popper 2017).
[52] See reference (Griffin and Shams 2018).
[53] See reference (Protos 2021).
[54] See https://news.coincu.com/101713-alameda-research-ftx-market-manipulation/ for an overview over this developing issue.

volume and the importance of Tether in various DeFi applications, a permanent depegging and subsequent crash would potentially lead to a collapse of the entire cryptocurrency space. Following the collapse of Terra/Luna, around 10 billion USDT where redeemed, about half of which was swapped by market makers to the less risky USDC.

Considering the issue of cryptocurrency valuation, the importance of the stability guarantee of stablecoins cannot be understated. The supposedly backed stablecoins effectively propagate their real value through direct trading and various decentralized finance applications throughout the entire crypto ecosystem. From a certain perspective, the current valuation of around 1 trillion dollars for the entire crypto market can be treated as a 6.66-fold leverage against the 150 billion US in stablecoins. Due to the intransparency of the current top stablecoin Tether[55], along with its numerous hints at insufficient assets for backing, this is another risk factor that is hard to quantify in the valuation of cryptocurrencies, but it is important to be recognized.

4.4 BTC-future price manipulation on the Binance exchange

Cryptocurrency exchanges differ from traditional regulated exchanges by unifying two important functions: the function of a brokerage, where buy and sell orders can be placed and which the exchange then executes, is being combined with the function of a clearinghouse, where the orders from brokerages are matched such that buyers and sellers indirectly interact. In traditional markets being on the downside of leveraged bets is handled by margin calls, which allow a time window to either post additional collateral or liquidate a position. The broker likewise keeps a margin with the clearinghouse, such that there is a two-stage buffering system in place to ensure the delivery and execution of all orders due to compartmentalization[56]. On crypto exchanges the unison of the brokerage and clearinghouse function can lead to a „failure-to-deliver“ situation, where the exchange owes more to the user than it can actually pay. This can be either be resolved by paying the difference from an

[55] In end-July 2022 several indications point at massive short positions against the stablecoin Tether, betting on its depegging from USD and a subsequent collapse of the entire crypto ecosystem.
[56] See reference (Ranger, An Anatomy of Bitcoin Price Manipulation 2022).

insurance fund, in which case the exchange loses money, or by auto-deleveraging the customer, such that the customer loses money[57].

Another critical element for the usage of excessive leverage on unregulated cryptocurrency exchanges is the issuance of perpetual future contracts, also known as perpetual swaps. The Binance BTCUSDT links the Bitcoin spot price directly to USDT Tether. Perpetual swaps do not have an actual delivery date and all differences between the spot price and future price are settled in regular intervals of 8 hours. Therefore, only relatively small amounts of money must be handed in each direction[58], which allows to push the leverage boundary as long as the settlements can be serviced. It has been shown that most of the price discovery and volatility of Bitcoin originates from the Binance exchange and its BTCUSDT perpetual futures contract[59].

In a time-series analysis of the full Binance orderbook at millisecond resolution it is shown by Ranger[56] how the Bitcoin price is manipulated during a short squeeze on July 26th, 2021, which is exemplary for multiple such occurrences throughout recent years. In a normal market the perpetual futures price and spot price move with minimal spread variance due to high-frequency trading bots that exploit any such potential arbitrage opportunities. Due to the placement and subsequent cancellation of large sell orders in a ladder pattern, an upwards price momentum is triggered due to the interaction with automated trading bots. When the market making bots that close the bid-ask spread on the perpetual futures side are suddenly turned off, the futures market almost instantly becomes illiquid whereas the spot market remains intact. Triggered by the upwards price movement, the liquidation of the most leveraged BTCUSDT positions in this environment leads to a sudden price

---

[57] The Binance exchange keeps an insurance fund specifically for dealing with the high level of leverage it allows for its traders. On May 19th, 2021, the USD/BTC price dropped by over 30% and USD/ETH by over 45% in a flash crash. Somewhat surprisingly, the Binance exchange trading platform itself crashed and was out for about 90 minutes. BTC and ETH prices hat recovered to pre-crash levels at that time, such that leveraged short positions where no longer in the money. The somewhat convenient outage of Binance may be explained by insufficient funds in the insurance fund, see reference (Alexander 2021).

[58] This is similar to contract-for-differences (CFD) trading in traditional markets, but uses standardized futures contracts instead of individualized CFD contracts.

[59] See references (Alexander, Heck and Kaeck, The Role of Binance in Bitcoin Volatility Transmission 2021) and (Alexander, Almost All Bitcoin Price Transmission Comes from Binance 2021).

spike of the futures contract, which is then transmitted by the arbitrage bots to the spot market, effectively driving up the price of Bitcoin.

Such price manipulations are the result of concentrated efforts by highly sophisticated market participants, allegedly by high-frequency cryptocurrency trading operations Alameda Research and Cumberland Capital. Any such market manipulation would be illegal in traditional markets and is much more difficult to execute due to the separation of functions. Due to the offshore nature of the quasi-unregulated Binance exchange and other questionable parties in the crypto space, such activities are much harder to identify and cannot be prosecuted. With regard to the valuation task, such price manipulation strategies are difficult to quantify but relevant to acknowledge.

4.5 Utility and commodity aspects of altcoins

The Ethereum monetary unit ETH differs from Bitcoin BTC in one crucial aspect: it is utilized as “gas” to process any transactions, execute smart contracts and run decentralized applications on the EVM, thus turning ETH from a pure store of value into a commodity powering the decentralized ecosystem. The EVM compute execution gas fee for smart contracts reflects an actual real-work expense for the provision of hardware and the energy spent, such that it represents natural economic incentives for the efficient usage of those limited resources. Therefore, as long as running decentralized applications has any value, ETH has an intrinsic value in addition to the pure store-of-value functionality of Bitcoin.

From the perspective of decentralized applications, specifically decentralized finance applications, ETH has numerous usage cases. For example, users can collateralize ETH through the MakerDAO project to back the stablecoin Dai, which earns them a variable interest. Decentralized applications like Aave, Compound, etc. allow users to stake (“lock up”) their assets to earn an interest, which the companies use for loans to other users, for example to take a leveraged bet on a short-term trading opportunity. Lending can also be used to for shorting the price development of cryptocurrencies, analogous to shorting stocks in traditional markets. This turns ETH into a yield-bearing asset, where the yield represents the risk taken on by the lender, as recently shown by the bankruptcy of the cryptocurrency lending company

Celsius in the aftermath of the Terra/Luna stablecoin collapse. Furthermore, there is the additional risk of impermanent loss if the trading value of a staked asset goes down while it is locked up, i.e., the owner cannot act on the this.

Likewise, ETH can be collateralized to enable swaps on decentralized exchanges by staking assets into liquidity pools for certain cryptocurrency pairs. For example, if one or many users stake their ETH and USDT into an ETH-USDT pool, this allows decentralized cryptocurrency exchanges like Uniswap or Sushiswap to provide the necessary liquidity for the service of swapping ETH and USDT[60]. The interest earned for the staking of those pools depends on the volatility and trading volume of the involved cryptocurrencies. A significant volume of those swap pools involves stablecoins, which links their real-world value into the entire crypto ecosystem, as mentioned before. Smart routing algorithms allow to exchange any cryptocurrency into any other (within the limits of the blockchain the DEX is operating on), potentially by utilizing multiple swap pools if no direct swap pool for the requested exchange exists or has insufficient liquidity.

The minting of new coins or tokens on a blockchain is usually accompanied by the creation of various swap pools pairing with stablecoins, which provide the exchange liquidity required for the trading of the new cryptocurrency. Due to the novelty, limited liquidity, and small trading volume, staking such coins usually earns particularly high interest rates, which quickly diminish once the cryptocurrency becomes more established. An entire industry called “yield farming” has been formed around the process of staking the highest interest-earning coins and quickly moving over to the next once the interest shrinks to more conservative levels—a process, that quite often completely ignores any counterparty risks associated with such novel coins. In this context, the interest yield on staked cryptocurrency becomes relevant as an indirect risk measure. It is typically specified in the annual percentage yield (APY), which reflects the real rate of return due to immediate compounding

[60] In practice, such swaps are executed as smart contracts on the blockchain the DEX is running on, using automated market maker (AMM) and routing algorithms. While the major stablecoins are natively provided on most large-scale blockchain projects, cryptocurrencies from other blockchains cannot be directly swapped. Instead, there are wrapper tokens to represent those foreign cryptos, e.g., WBTC as “wrapped Bitcoin” on the Ethereum chain. Modern blockchain 3.0 projects like Cosmos or Polkadot aim to build an ecosystem supporting a better blockchain interoperability.

of interest $\text{APY} = \left(1 + \frac{r}{n}\right)^n - 1$ where r is the nominal interest rate and n the number of compounding periods per year.

Another application of decentralized finance are so-called oracles, which are real-world data providers for usage in smart contracts. For example, an insurance smart contract may lock up a certain amount of crypto assets to be paid to the customer if his house burns down. An associated oracle, i.e., a person or company in the real world, would keep an eye on the insured house and provide the necessary information to the smart contract if the house indeed burns down. To avoid misuse or fraud, an oracle must stake a collateral related to the value of the provided information involved in the smart contract to ensure its trustworthiness. If the information turns out to be fraudulent, the collateral is seized, otherwise the oracle earns an interest for proving the real-world information. Such tie-ins to the real world, which can be used in transparent and tamper-proof automated payment processes, provide an obvious value to the crypto space and therefore to the involved cryptocurrencies.

Ethereum and other cryptocurrencies that support general-purpose applications therefore become drivers of an entire financial ecosystem with real-world ties and applications. Those cryptocurrencies serve as stores of value (like Bitcoin), but in addition work as commodities empowering the respective blockchain's contract execution and act as interest-earning assets.

To measure and quantify the utility aspect of these altcoins, the total amount of staked cryptocurrency can be used, since staking is the primary means of enabling DeFi applications. This is referred to as Total Value Locked (TVL). Overall, the DeFi utility aspects of a cryptocurrency are quantifiable dimensions and provide a suitable additional input for any cryptocurrency valuation model, where they are applicable due to the general-purpose programmability via smart contracts.

5 Short-term ETH valuation forecast using DeFi TVL metric

The complexity of the cryptocurrency valuation task has been shown in the review of prior attempts (chapter 3) and the discussion of qualitative problems (chapter 4). Based on those considerations, a quantitative approach for short-term price predictions of smart-contract-enabled cryptocurrencies is presented and backtested against data. This work is focused on the Ethereum blockchain as the largest representative of this group.

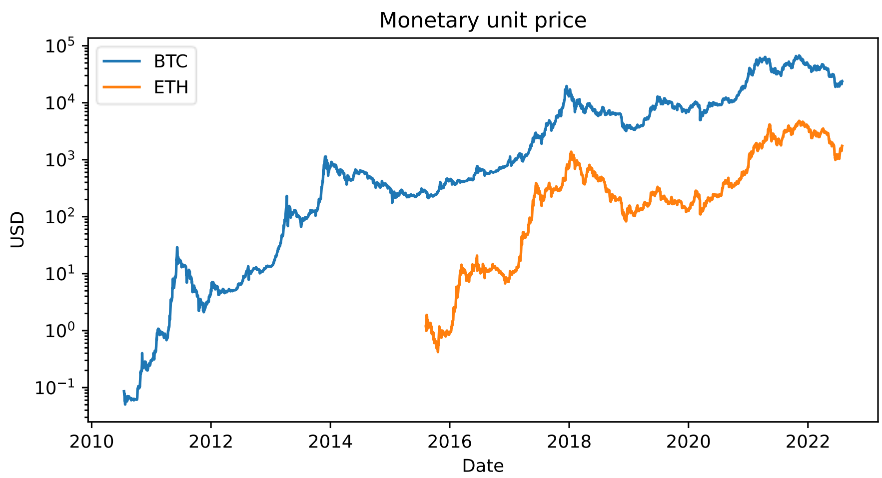

Figure 6: Bitcoin BTC and Ethereum ETH price in US dollar
(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

Specifically, a cross-correlation analysis of historical time series is performed on BTC and ETH market capitalization data and the TVL metric of the Ethereum blockchain. Motivated by the results, several machine learning models for next-day predictions of the ETH price are trained on different combinations of those time series as an input. A performance analysis shows that the ETH next-day-price can be reasonably well predicted from non-ETH historical data once the TVL metric is included.

5.1 Correlations between cryptocurrencies and TVL

Individual cryptocurrency price levels and the total cryptocurrency market capitalizations are often empirically observed to be moving in tandem. Comparing the BTC and ETH price over time in Figure 6, peak pricing levels are hit almost at the same time, e.g., the 2017/18 BTC and ETH peak. Figure 7 shows the relative price behavior, i.e., the ETH price in terms of BTC units. After the early years of

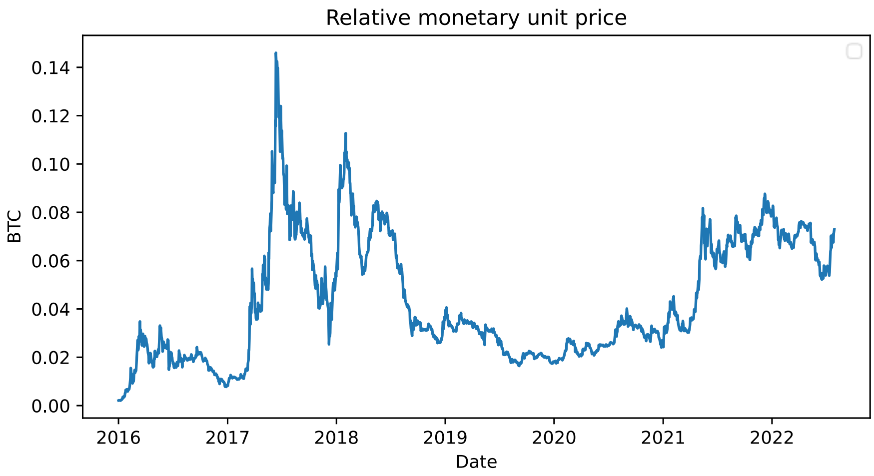

Figure 7: Ethereum ETH price in Bitcoin BTC

(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

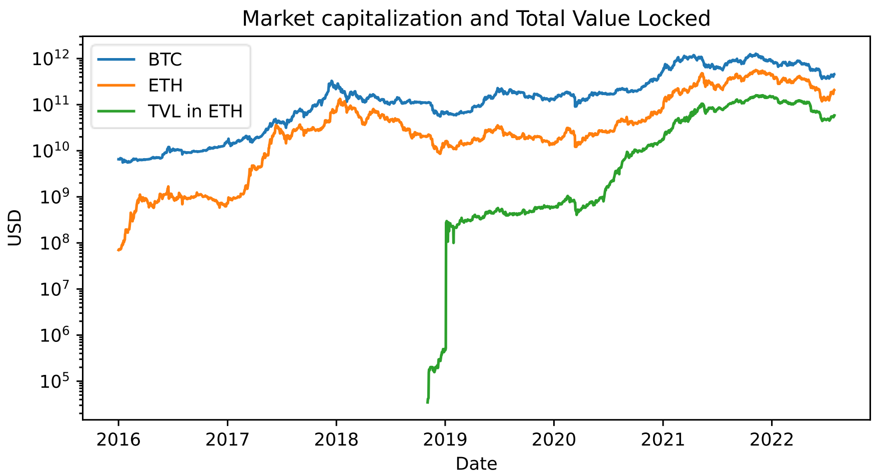

Figure 8: BTC and ETH market cap. and TVL in Ethereum

(Created using data retrieved from CoinMetrics and Defi Llama,
see https://coinmetrics.io and https://defillama.com.)

Ethereum and passing through the high volatility of the 2017/18 crypto bubble, the relative price stabilizes beyond the end of 2018. This coincides with the introduction of decentralized finance, compare Figure 8. The sharp rise in relative pricing coincides with the mid-2021 crypto price rise and the rise of decentralized finance. Even after the 70% retraction of the crypto market in early-2022 the relative ETH/BTC price level remains higher compared to the 2019 to 2021 timeframe. This hints at a gradual shift of Ethereum from a pure store-of-value digital asset towards at least a partial utility character, as discussed in section 4.5.

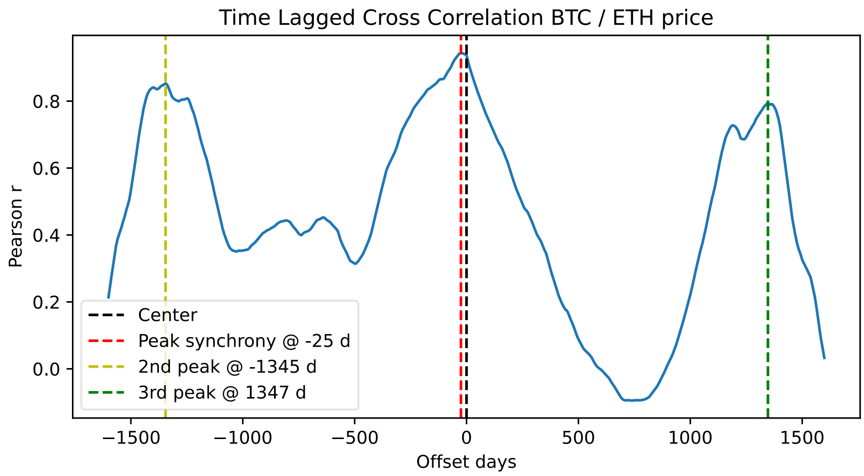

Figure 9: TLCC between Bitcoin and Ethereum price

(Created using data retrieved from CoinMetrics, see https://coinmetrics.io.)

The logarithmic price graph in Figure 6 hints at a high correlation between the pricing behavior of the two different cryptocurrencies. The naïve Pearson correlation for two distributions $X = \{x_i\}$ and $Y = \{y_i\}$ is defined by

$$\mathrm{Corr}(X,Y) = \frac{\mathrm{Cov}(X,Y)}{\sigma_X \sigma_Y} = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^{n}(x_i - \bar{x})^2 \cdot \sum_{i=1}^{n}(y_i - \bar{y})^2}$$

where $\bar{x}$ and $\bar{y}$ are the centers of the two distributions. This standard metric shows a high correlation of $r = 0.92$ when applied to the BTC and ETH price starting from January 1st, 2016, which removes the boundary artifacts of the first months of ETH.

A more sophisticated analysis of the continuous pricing time series can be done using the time-lagged cross-correlation (TLCC), where essentially one time series is shifted relative to the other and for each shift the Pearson correlation is computed. The shift offset of peaks of such a TLCC computation indicate the relative offset of time-series, i.e., at which offset both are most highly correlated. The resulting graph

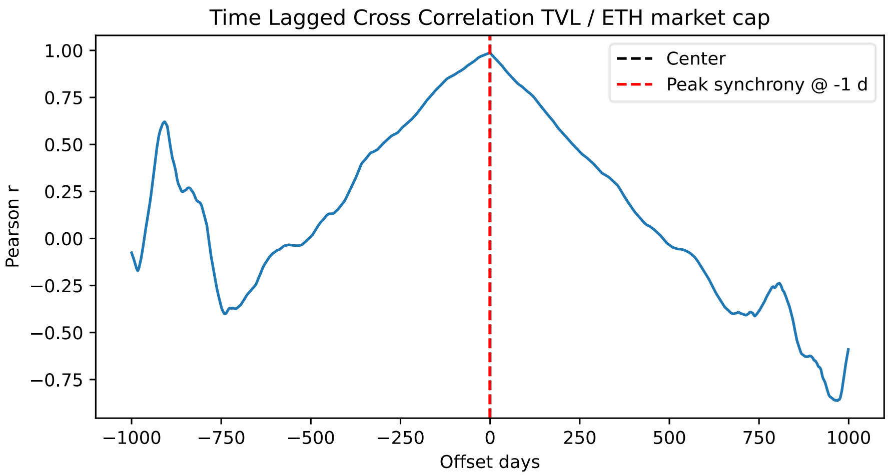

Figure 10: TLCC between Ethereum market cap. and TVL in Ethereum

(Created using data retrieved from CoinMetrics and Defi Llama,
see https://coinmetrics.io and https://defillama.com.)

of time-shifted Pearson correlations is shown Figure 9. For the BTC and ETH pricing data this TLCC analysis reveals a peak correlation[61] of $r = 0.94$ at an offset of 25 days.

While such a correlation is never an indication of causation, due to the similar nature of both digital assets a link in the pricing level can be assumed as a first approach. Based on this, one of the most naïve models to value ETH is therefore given by a time shifted BTC pricing model. In terms of an absolute valuation model the stock-to-flow model discussed in section 3.4 remains the best option, despite the various problems. An improved BTC-based ETH pricing model is being introduced in the next section.

A second observation of a similar kind shows a strong correlation between the Ethereum price and the TVL amount of decentralized finance applications on the blockchain, especially from mid-2020, see Figure 8. Performing the same TLCC analysis on the ETH market capitalization and the TVL on the Ethereum blockchain indicates an almost synchronous movement pattern between the time series with a peak Pearson correlation of $r = 0.987$, compare Figure 10.

[61] The 2nd and 3rd peak indicated in Figure 9 are partially border effects due to the relative time series shifts and coincide in part with the 3.5-year offset between the December 2017 and April 2021 bull run, i.e., the rough pricing shape of the two crypto peaks is somewhat similar. The entire considered data consists of 2402 days.

5.2 CNN-LSTM next-day prediction model description

Based on the strong correlation between the TVL and ETH market capitalization, as well as the BTC and ETH price correlation, a next-day-prediction model is used to analyze where most of the pricing information (and therefore the associated value information of the Ethereum cryptocurrency) can be recovered from. This is a rather pragmatic quantitative approach to identify and rank price driving metrics, but it does not give rise to a fundamental valuation approach. Various methods are available for time series predictions and forecasting, all of which have inherent tradeoffs. In this work the relatively recent combination of convolutional and long-short-term memory neural networks is being used, called an CNN-LSTM model, which has been successfully applied to stock price forecasting[62].

Convolutional neural networks (CNNs) are deep neural networks[63], that are often applied in the context of image classification (2-dimensional) and time series analysis (1-dimensional). The features of the input data to the network are extracted using multiple convolution filters, which are automatically determined during the training process of the neural network. This allows to design a model without prior knowledge regarding the important aspects of the input data and eliminates the need for hard feature engineering using external knowledge. For example, given enough training data, a pure CNN model could be applied to perform a form of technical analysis on fixed-length time series slices without prior specification of the various trading patterns commonly identified. This aspect is in part utilized here, however, the amount of training data available is rather restricted.

Long-short-term-memory models (LSTMs) are deep neural networks, as well, specifically a type of recurrent neural networks (RNNs) which have feedback loops. Using artificial neurons as gates to activate or deactivate functionalities like writing, retrieving, or clearing of data, LSTM units behave like a computer memory cell. This allows to store historical state information during operation, which makes LSTMs particularly suited to deal with complex sequential or time series data.

---

[62] See reference (Lu, et al. 2020).

[63] A deep neural network is an artificial neural network with multiple hidden internal layers of connected neurons.

By combining both types of neural networks, the LSTM part of the network therefore captures the time-like behavior of the distribution of local features that have been identified from the various CNN filters. The data considered here are time series of the daily closing prices for BTC and ETH as well as the daily TVL amounts on the Ethereum blockchain. This data is organized in fixed-length segments as described in the next section.

The CNN portion of the model uses 64 one-dimensional filters[64] of size 5, followed by a dropout layer with a 20% chance of zeroing the output value to avoid overfitting[65] during training. The 64 output feature maps resulting from the CNN operation are then fed into a 20-unit LSTM layer, meaning that it consists of 20 internally connected LSTM cells capable of reproducing sequences of up to 20 elements. Various other combinations of the network parameters chosen here are valid and can be argued for, but a systematic sampling of this vast hyperparameter space is outside the scope of this work. Finally, the output of the LSTM layer is fed into a densely connected network layer that culminates into a single output value, which is the predicted ETH price of the day following the input time slice.

5.3 Data preparation and model training

The input data to the model is a combination of the daily closing prices for Bitcoin and Ethereum[66] as well as the daily TVL amount[67] on the Ethereum blockchain. All pricing and TVL time series are linearly rescaled to a $[0,1]$ interval to improve training efficiency and numeric stability[68]. In the following the data ranging from March 1st, 2019 till July 30th, 2022 is considered, encompassing 1248 days of data. Only the data up to December 31st, 2021 is used during training and the rest is kept

[64] Non-systematic experiments during the development stage of the model have indicated, that a reduction of the number of filters to 32 yields a much worse result for the predictive power. The chosen number of 64 CNN filter maps empirically leads to usable results. Due to the issue of over-fitting and data scarcity, this choice follows the principle to use as few parameters as necessary and as little as possible.

[65] Overfitting is a general problem in machine learning approaches, where a model learns the training data with minimal error but fails to generalize to new and previously unseen data. By randomly introducing "errors" to the internal computation via a dropout layer this issue can be reduced.

[66] Data retrieved from the free tier of CoinMetrics, see https://coinmetrics.io.

[67] Data retrieved from Defi Llama, see https://defillama.com.

[68] Data used in machine learning applications is typically rescaled to the same order of magnitude and for simplicity reasons to either the $[-1,1]$ or $[0,1]$ intervals, depending on applications. This scaling is related to vanishing gradients of non-linear functions during training, which lead to numeric instabilities and/or extremely slow training progression.

for a validation of the training, i.e., to test whether or not the model generalized its predictive powers to previously unseen data. The data is being organized into sliding window samples of 21 sequential days on the input side. Therefore, the model is making next-day ETH price predictions based on the daily data of the prior three weeks.

The untrained network starts with a semi-random initialization[69]. An input sample from the training data is fed into the network and the internal computations throughout all layers are performed while keeping their respective intermediate values. At the output note of the network the computed value is compared to the expected output value and the difference is the current training error. Based on this error a so-called loss function is defined, typically based on the mean-squared-error. The goal of training a neural network—or machine learning model in general—is to minimize the loss function and therefore the error between the predictions computed by the model and the expected output value. Training a neural network is performed via the backpropagation algorithm. Layer by layer the computed error at the output stage is being split up and distributed as partial errors on the nodes of the prior layer based on the current state of the neural network. Those partial errors together with the internal temporary values is then used to modify the current state of the network with the goal of minimizing the error on this training sample. This process is repeated numerous times, either based on the number of training iterations or a certain target threshold on the loss function.

During batch training multiple input samples are send through the current state of the network, but the error computation and subsequent backpropagation of errors is only performed once the entire batch has been computed, i.e., the network is updated on batches of training data. This speeds up the training process considerably. Going through the entire set of all the training samples is called an epoch of the training process. The training of the CNN-LSTM model considered here is carried out in 200 epochs with a batch size of 8. Like for the choices of the neural network architecture and the hyperparameters, other values for the training process could have been considered but a full-scale analysis is outside the scope of this work. The

[69] For numeric purposes and to mitigate the vanishing gradient problem, the neural network parameters are typically initialized with random values not to close to zero, thus semi-random.

actual training is carried out using Google Tensorflow[70], an industry-standard software for complex machine learning tasks involving neural networks.

The goal of the CNN-LSTM model is to predict the next-day price of the Ethereum monetary unit based on sequential time series data from the prior 21 days. Based on the model architecture, data preparation and training parameters outlined here, three different models are trained on different input data sets.

5.4 Model 1: ETH price prediction using BTC price

The first model is trained on Bitcoin price data. Each input consists of a 21-day time slice of BTC prices to predict the ETH price of the day following the end of the

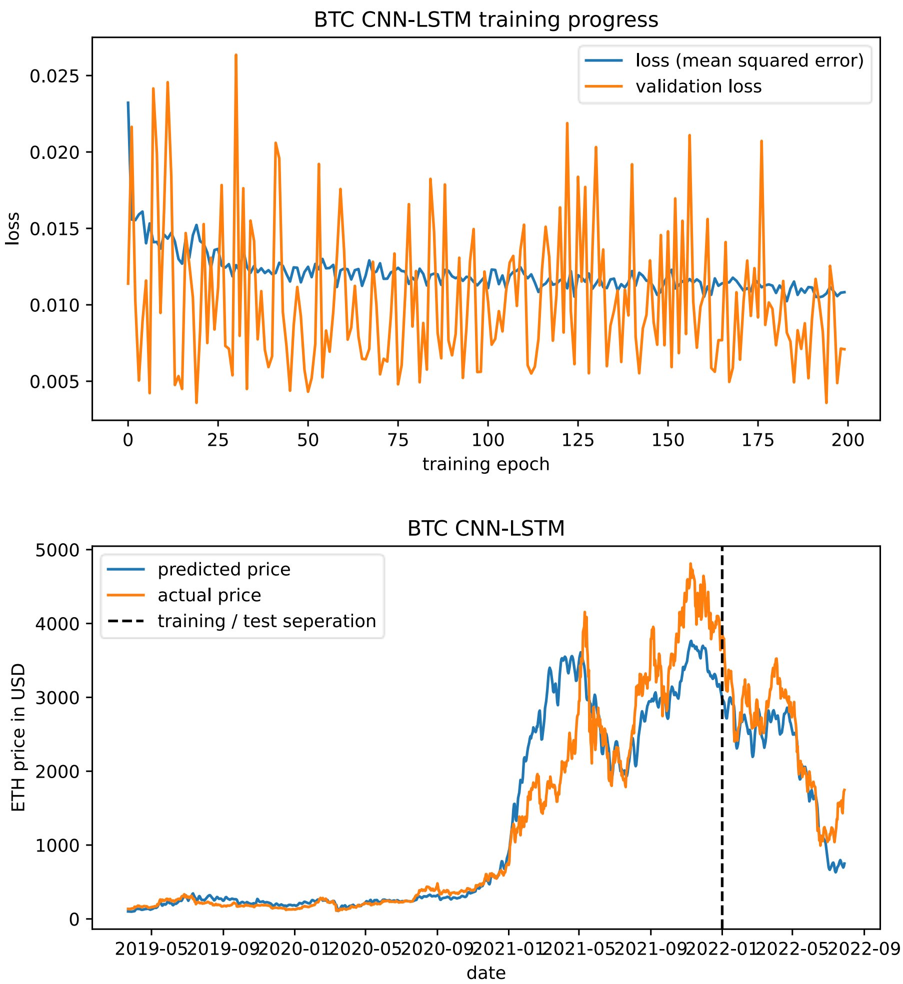

Figure 11: BTC-based model training progress and prediction

[70] See https://www.tensorflow.org for details. Alternative tools like Facebook PyTorch, Matlab, etc. could have been used instead.

time slice. The input CNN layer of the model has 384 trainable parameters, followed by 6800 trainable parameters in the LSTM layer and 20 trainable parameters in the densely connected last layer that brings all results together, leading to a total of 7205 free parameters of the model. Considering that only 1015 training samples of 21 days each are fed into the network (21315 numbers in total), this is a data-poor training situation. Therefore, the relative high number of 200 training epochs together with the usage of dropout layers both at the CNN and LSTM level has been chosen to reduce the problem of overfitting the model to the training data.

The training progress and predictions of this BTC price-based model are shown in Figure 11. As can be seen the training error (loss) does not reduce significantly and remains relatively high even after several training epochs. More importantly, the validation loss keeps a very high level of fluctuations even at late stages of the training process. When inspecting the model prediction next to the actual ETH price, the significant deviations are obvious. While the Bitcoin price may capture the rough pricing behavior of Ethereum, both cryptocurrencies have different pricing nuances.

5.5 Model 2: ETH price prediction using TVL amount on Ethereum blockchain

The second model is trained on the TVL amount on the Ethereum blockchain. Each input consists of a 21-day time slice of TVL data, such that the considerations regarding the number of parameters and input data are analogous to the first model. Basically, the neural network is simply trained on different input numbers.

The results of the TVL-based model predictions and the training progress can be seen in Figure 12. The training error reduces to a much lower level compared to the BTC-based model, indicating that the TVL data is better suited to reproduce the training data set. This is not unexpected considering the strong correlation with minimal time-shift between both time series, which was identified in section 5.1. On the validation data set the error keeps wildly fluctuating and does not appear to

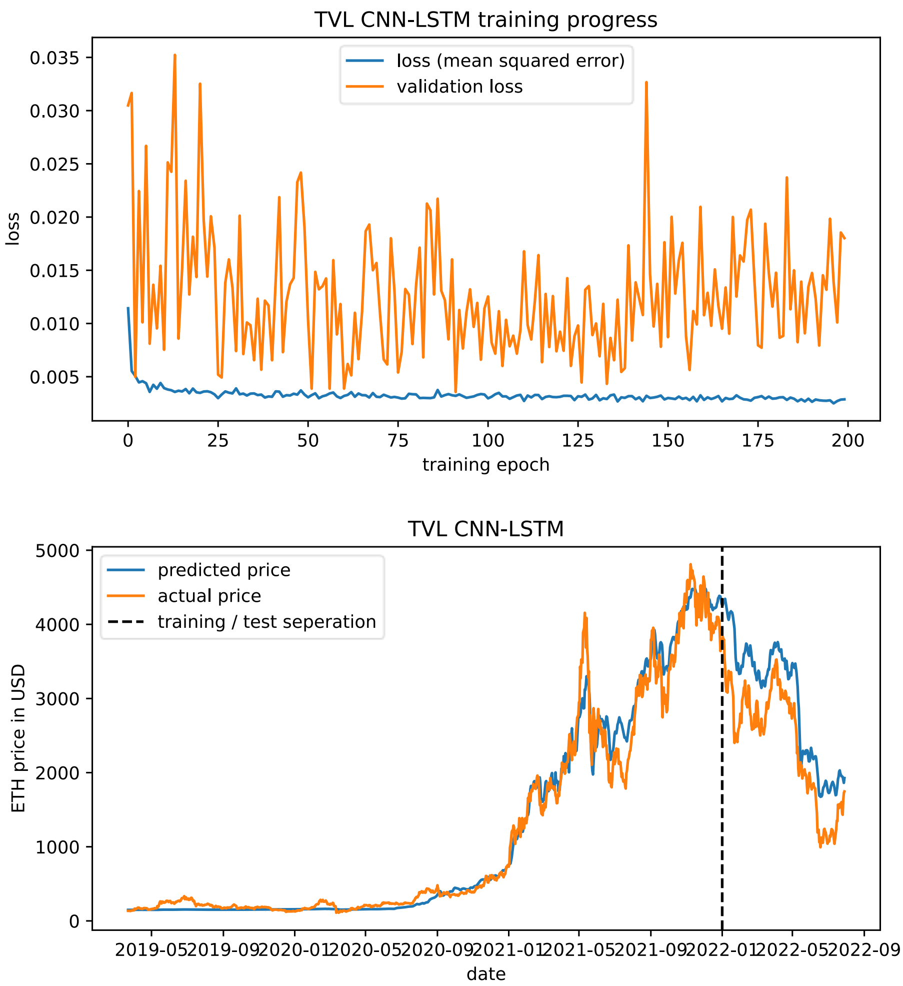

Figure 12: TVL-based model training progress and prediction

reduce significantly with more training epochs, even though the validation loss oscillates at a smaller magnitude compared to the BTC-based model. The TVL-based model therefore is an improvement compared to the BTC-based model on the training data set but does not appear to generalize particularly well.

5.6 Model 3: ETH price prediction using both BTC price and TVL on Ethereum blockchain

For the third model a combination of both the BTC price data and the TVL amount locked on the Ethereum blockchain is considered. Each input consists of two 21-day time slices of the respective data sets, which are fed as different channels[71] into

[71] An input channel of a CNN is a generalization of the concept of colors. In 2d image analysis applications images are delivered as a 3-channel input to the CNN, where the channels represent the colors.

the CNN part of the network. Thus, the input vector is 42-dimensional and the additional input dimensions lead to an increase of 704 trainable parameters at the CNN input layer due to handling two time series channels in the convolution, but the CNN output feature maps remain unchanged. The LSTM layer remains unaffected, such that the total number of trainable parameters increases to 7525, a 4.4% increment compared to the prior models. Regarding the available input data, the number of samples remains the same, but 42730 numbers are effectively fed into the network training. This constitutes still a rather data-poor training situation, however, an improvement compared to the other two models.

The training progress and comparison between the predicted and actual ETH price can be seen in Figure 13. The significant improvement compared to the prior models is obvious: while the error on the training set is improved, the validation loss shows a significant reduction in oscillation during the progression of training

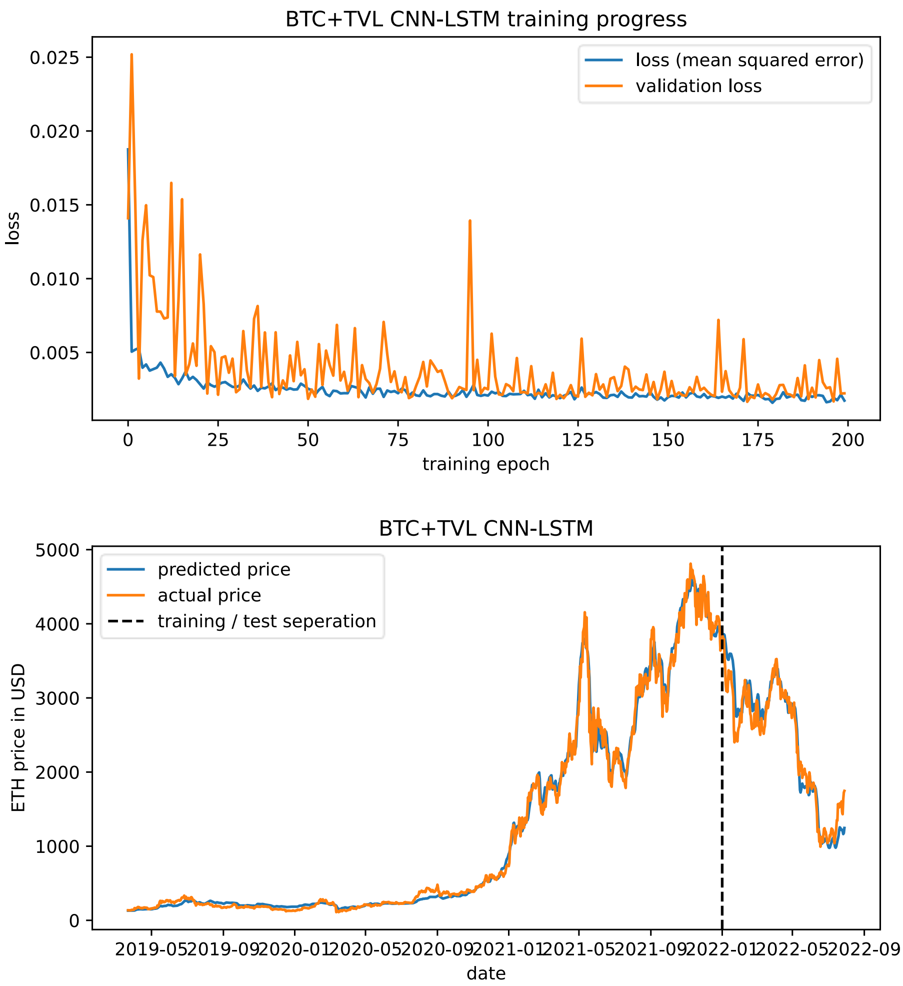

Figure 13: BTC+TVL-based model training progress and prediction

epochs and settles at the same level as the training loss. This indicates that the model is much better able to generalize to previously unseen data, which can be qualitatively verified in the close movement of the predicted and actual ETH price beyond the training data threshold. While isolated BTC prices and TVL amounts produced rather weak predictive results, a combination of both information yields a reasonably well performing next-day price predictor, even on unseen data.

5.7 Reference Model: ETH price prediction using BTC, TVL and ETH

To benchmark the results of the prior three models, a reference model is trained on all the available data: three time series slices of the past 21-days of BTC, TVL and ETH are served as input for the training. The CNN input stage is expanded analogous to the previously discussed BTC+TVL combined model. This increases the number of trainable model parameters in the CNN input layer to 1024, such that

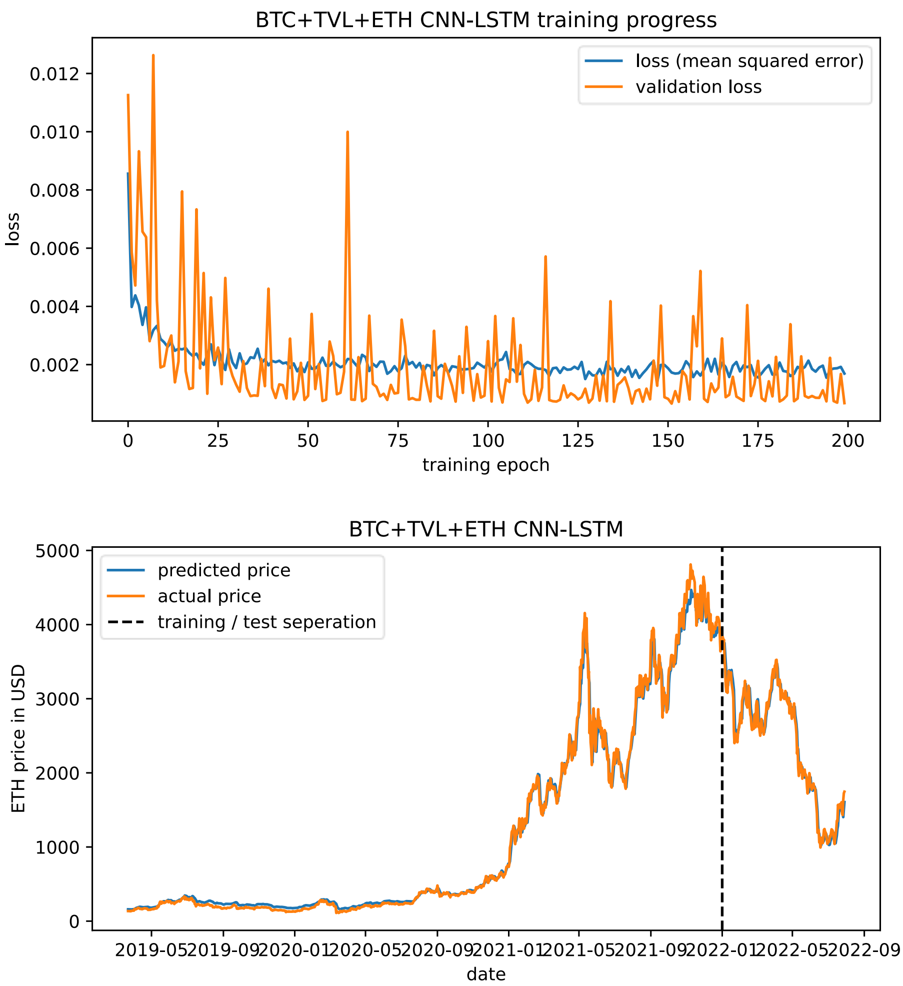

Figure 14: BTC+ETH+TVL-based reference model training progress and prediction

the reference model has 7845 trainable parameters in total. The number of training samples remains unchanged, which leads to 63945 numbers effective being used during the training of the model.

The results of the training process can be seen in Figure 14. With the exception of the price peak reached in late-2021, the predicted and actual ETH price levels remain very close. The improved performance can be seen in the graph of the training loss, which reaches lower levels compared to the other models.

5.8 Analysis of model performance

The three models trained in the previous sections indicate that it is possible to predict the short-term price behavior of the Ethereum monetary unit based on the Bitcoin price and the DeFi metric TVL. Furthermore, the prediction quality worsens significantly if only one of the two input time series is being used. Therefore, both the BTC price and TVL amount add substantial information about the behavior of the ETH price. Based on the training progress and the final loss of 0.003 for the TVL-based model compared to 0.01 for the BTC-based variant after training, the TVL metric captures significantly more information about the ETH price, compare Figure 11 and Figure 12.

Another aspect to consider is the relative novelty of entire DeFi ecosystem in early-2019, which can be seen in the sharp rise of the TVL value in late-2018/early-2019, compare Figure 8. The training of the model was therefore restricted to dates after March 1st, 2019, to remove exceptional boundary effects. However, the TVL amount was still almost two orders of magnitude smaller compared to the market capitalization of Ethereum. This gap did significantly reduce in the subsequent three years.

Considering the results of the correlation analysis in section 5.1, it appears that the Bitcoin price behavior front-runs the Ethereum price when considering larger price changes, whereas the finer details of the price behavior are better captured in the total value locked in DeFi applications on the blockchain. Such short-term/long-term dependencies can be analyzed using spectral decompositions, which can be used to identify frequency-based dependencies, but this is outside the scope of this work. Such decompositions are also the basis for alternative prediction models, e.g., based on SARIMA models.

6 Conclusions

This work indicated that the DeFi metric TVL captures a significant portion of the short-term price behavior of the Ethereum cryptocurrency. Based on the relative predictive power of the CNN-LSTM models considered here, the TVL metric is a stronger determinant of the ETH price than the largest cryptocurrency Bitcoin, which otherwise still dominates in terms of market capitalization.

As reviewed in chapter 3, numerous long-term valuation forecasting approaches for Bitcoin have so far not settled the issue of finding fundamental metrics akin to established valuation models for traditional assets, e.g., discounting cashflows for companies. The stock-to-flow model so far has managed to provide the best macro-scale success by empirically capturing the sharp order-of-magnitude price rises associated with the miner reward halvings in the Bitcoin ecosystem, albeit lacking a sound statistical foundation. Moreover, this forecasting approach depends on the hard-coded intervals of those halving events and does not generalize to other cryptocurrencies. Specifically in the case of Ethereum, where the monetary policy has changed on multiple occasions and a partial burning of gas fees utilized on the EVM was introduced, accurate predictions based on ETH supply and issuance are sketchy at best. Moreover, the currently ongoing transition from a proof-of-work to a proof-of-stake consensus mechanism entirely changes the implied economic incentives underlying the proof-of-work model which may have significant effects on valuation. This also invalidates any modelling based on costs of production. Measuring the network throughput or number of transactions on the blockchains has also become rather problematic. Major blockchains have reached technological throughput limits and significant portions of the transactional volume has been outsourced to 2nd-layer extensions, like the Lightning network for Bitcoin or various EVM-compatible chains anchored to Ethereum, e.g., Polygon. Ultimately, this overview of the state of quantitative valuation approaches shows the lack of generally accepted fundamental approaches and the unfeasibility to reuse some Bitcoin-specific attempts to smart-contract-enabled blockchains.

The conceptual, legal, and fraudulent issues surrounding the crypto space described in chapter 4 add to the complexity of the valuation problem. One of the key features associated with cryptocurrencies—to provide a censorship-free payment system without the possibility of governmental intervention—breaks down in real-world scenarios for most users. Ironically, criminals seem to have a better grasp on those limitations and understand to use them to their advantage. Various services and blockchains have been developed, sometimes by questionable parties, with the goals of anonymity and transaction obfuscation in mind. This has turned cryptocurrency in the default payment method for illegal activities of all kinds.

However, the true ticking time bomb is the highly dubious situation of stablecoins, which are the backbone of the entire DeFi ecosystem on one hand, but are unregulated, unaudited, and most likely at least partially unbacked on the other hand. The entanglement of some of the largest (and suspiciously unregulated) players in the crypto space—stablecoin issuers, exchanges as well as sophisticated high-frequency traders—has created a highly fragile house of cards for any serious investor or short-term speculator.

Due to those long-term uncertainties and a lack of fundamental metrics, the model presented in this work is focused on short-term price predictions as a valuation proxy. To identify the key driver of the 24-hour price change in ETH, multiple CNN-LSTM-based neural networks were trained on 21-day segments of BTC and TVL historical data. This revealed an improvement of the prediction quality if TVL data instead of BTC prices were used, indicating that a significant portion of the short-term future price information is contained in the TVL DeFi metric. Even better predictions could be recovered from simultaneous usage of BTC and TVL time series segments. The Ethereum value is therefore linked both to the store-of-value aspects of Bitcoin and the utility aspects of decentralized finance—with a tilt to the latter.

There are various avenues for an improvement of those findings: the main problem of the results presented here is the relative data scarcity both for training and validation of the models. Using a higher-frequency resolution of the data could significantly improve (or weaken) the quality of the results, i.e., strengthen or weaken the predictive link identified between the TVL metric and the short-term ETH price. This requires only minor modifications to the base model to cope with the different input and output structure.

Higher-resolution data could also be used in the correlation analysis, which originally motivated the selected approach to perform short-term predictions of the ETH price based on the BTC price and TVL metric. The time-lagged Pearson cross-correlation used in section 5.1 implicitly assumes a fixed relative synchroneity of both time series. Instead, an initial 7-day delayed trend following at the beginning of the time series could gradually reduce to a 1-day delay at the end, for example. Such situations can be analyzed using dynamic time warping during the matching. As an alternative, instead on considering the global correlation of the entire time-shifted time series, windowed local portions can be considered, which provides a more fine-grained view on the matching and mismatching segments of the time series. For example, the identified 25-day time delay between the BTC and ETH price levels presumable could have been reduced in more recent time frames.

The short-term prediction model itself can be improved on various aspects. A systematic analysis of the various hyperparameters chosen in this work could identify a potentially much better model setup. For example, the number of CNN filters at the input stage of the model could be increased or reduced, as well as the size of the actual filters. The LSTM stage likewise can be adjusted in the number of connected units, thus changing the internal temporal dependencies. In general, techniques from the field of AutoML could be applied to automate the search for optimal hyperparameters of the model architecture.

Regarding the valuation task of altcoins, the next step would be to apply the same model to other smart-contract-enabled blockchains and their respective TVL DeFi metric. If the same relationship can be established for other coins, more general valuation models involving multiple TVL metrics from different blockchains as input could be investigated—either individually or in various combinations.

Ultimately, the link between the value locked in decentralized finance applications and the price of the underlying utility token of the blockchain, which is suggested by the findings of this work, provides a starting point for future research in the valuation of cryptocurrencies with general-purpose programmability.

Appendix

A The top 10 cryptocurrencies

Name	USD Price	Market Capi-talization	24h Trading Volume	Circulating Supply
Bitcoin	$21,438.50	$408,191M	$25,523M	$19M BTC
Ethereum	$1,487.69	$179,426M	$17,016M	$122M ETH
Tether	$1.00	$65,858M	$44,840M	$66B USDT
USD Coin	$1.00	$55,125M	$6,442M	$55B USDC
BNB	$258.73	$41,370M	$967M	$161M BNB
Binance USD	$1.00	$17,758M	$5,277M	$18B BUSD
XRP	$0.3393	$16,283M	$981M	$48B XRP
Cardano	$0.4743	$15,834M	$635M	$34B ADA
Solana	$37.28	$12,774M	$1,189M	$356M SOL
Dogecoin	$0.06313	$8,330M	$379M	$133B DOGE

Table 1: Market information of top 10 cryptocurrencies on July 27th, 2022

References

Alabi, Ken. "Digital blockchain networks appear to be following Metcalfe's Law." *Electronic Commerce Research and Applications* 24 (2017): 23-29.

Alexander, Carol. "Almost All Bitcoin Price Transmission Comes from Binance." 2021. https://www.coalexander.com/post/almost-all-bitcoin-price-transmission-comes-from-binance.

—. "Binance's Insurance Fund." August 3, 2021. https://www.coalexander.com/post/binance-s-insurance-fund.

Alexander, Carol, Daniel Heck, and Andreas Kaeck. "The Role of Binance in Bitcoin Volatility Transmission." *arXiv.* 2021. https://doi.org/10.48550/arXiv.2107.00298.

Ammous, Seifedean. *The Bitcoin Standard.* Wiley John & Sons, 2018.

Andres, Sebastian. *Ethereum in a Nutshell.* Independently published, 2021.

Antonopoulos, Andreas M. *Mastering Bitcoin.* 2nd. O'Reilly UK Ltd., 2017.

Antonopoulos, Andreas M., and Wood Gavin. *Mastering Ethereum.* O'Reilly UK Ltd., 2018.

Antonopoulos, Andreas M., Olaoluwa Osuntokun, and René Pickhardt. *Mastering the Lightning Network.* O'Reilly UK Ltd., 2021.

Aspembitova, Ayana T., Link Feng, and Lock Y. Chew. "Behavioral structure of users in cryptocurrency market." *PLoS One* 16, no. 1 (2021).

Bier, Jonathan. *The Blocksize War.* Independently published, 2021.

"Bitcoin Energy Consumption Index." *Digiconomist.* 2022. https://digiconomist.net/bitcoin-energy-consumption/.

Booth, Jeff. *The Price of Tomorrow.* Stanley Press, 2020.

Briscoe, Bob, Andrew Odlyzko, and Benjamin Tilly. "Metcalfe's law is wrong." *IEEE Spectrum* 43, no. 7 (July 2006): 34-39.

Burger, Marcel. "Challenging PlanB: a review of modelling Bitcoins value with scarcity." *Medium.* 2019. https://medium.com/amdax-asset-management/challenging-plan-b-a-review-of-modelling-bitcoins-value-with-scarcity-3d4e7e141286.

—. "Reviewing "Modelling Bitcoin's Value with Scarcity" — Part III: The Fall Of Cointegration." *Medium.* 2020. https://medium.com/amdax-asset-management/reviewing-modelling-bitcoins-value-with-scarcity-part-iii-the-fall-of-cointegration-ec5a8267098a.

—. "Reviewing "Modelling Bitcoin's Value with Scarcity" —Part II: The hunt for cointegration." *Medium.* 2019. https://medium.com/amdax-asset-management/reviewing-modelling-bitcoins-value-with-scarcity-part-ii-the-hunt-for-cointegration-66a8dcedd7ef.

Buterin, Vitalik. "Ethereum Whitepaper." 2014. https://ethereum.org/669c9e2e2027310b6b3cdce6e1c52962/Ethereum_Whitepaper_-_Buterin_2014.pdf.

Cipolaro, Greg, and Ross Stevens. "The Power of Bitcoin's Network Effect." *NYDIG.* November 2020. https://nydig.com/research/the-power-of-bitcoins-network-effect.

Diedrich, Henning. *Ethereum.* CreateSpace Independent Publishing Platform, 2016.

Dwork, Cynthia, and Moni Naor. "Pricing via Processing or Combatting Junk Mail." *Advances in Cryptology — CRYPTO' 92.* 1993. 139–147.

Europol. *Cryptocurrencies: Tracing the Evolution of Criminal Finances.* Europol Spotlight Report series, Luxembourg: Publications Office of the European Union, 2021.

Ferguson, Niels, Bruce Schneier, and Tadayoshi Kohno. *Cryptography Engineering.* Wiley, 2010.

Franco, Pedro. *Understanding Bitcoin.* Wiley, 2014.

Griffin, John M., and Amin Shams. "Is Bitcoin Really Un-Tethered?" 2018. https://dx.doi.org/10.2139/ssrn.3195066.

Harvey, Campbell R., Ashwin Ramachandran, and Joey Santoro. *DeFi and the Future of Finance.* Wiley, 2021.

Hayes, Adam. "A Cost of Production Model for Bitcoin." 2015. https://dx.doi.org/10.2139/ssrn.2580904.

Ige, Lanre. "Cryptoasset valuation techniques." 2018. https://s3.us-east-2.amazonaws.com/lanreige/Cryptoasset+Valuation+Techniques.pdf.

In Gold We Trust Report. *The stock-to-flow ratio as the most significant reason for gold's monetary importance.* 2012. https://ingoldwetrust.report/the-stock-to-flow-ratio-as-the-most-significant-reason-for-golds-monetary-importance/?lang=en.

Kripfganz, Sebastian. "The stock-to-flow model of Bitcoin: an introduction and critical assessment." *YouTube.* May 12, 2020. https://www.youtube.com/watch?v=u_VkwA5fkLA.

Lantz, Lorne, and Daniel Cawrey. *Mastering Blockchain.* O'Reilly UK Ltd., 2020.

Lu, Wenjie, Jiazheng Li, Yitan Li, Aijun Sun, and Jingyang Wang. "A CNN-LSTM-Based Model to Forecast Stock Prices." *Complexity* 2020 (2020).

Madureira, António, Frank den Hartog, Harry Bouwman, and Nico Baken. "Empirical validation of Metcalfe's law." *Information Economics and Policy* 25, no. 4 (December 2013): 246-256.

McKenzie, Patrick. "Tether Required Recapitalization In May 2022." *Kalzumeus.* 2022. https://www.kalzumeus.com/2022/05/20/tether-required-recapitalization/.

Menger, Carl. "On the Origins of Money." *The Economic Journal* (Oxford University Press) 2, no. 6 (1892): 239-255.

Metcalfe, Robert. "Metcalfe's Law after 40 Years of Ethernet." *IEEE Computer* 46, no. 12 (2013): 26-31.

Möser, Malte, et al. "An Empirical Analysis of Traceability in the Monero Blockchain." *Proceedings on Privacy Enhancing Technologies Symposium* 3 (2018): 143-163.

Nakamoto, Satoshi. "Bitcoin: A Peer-to-Peer Electronic Cash System." 2008. https://bitcoin.org/bitcoin.pdf.

Peterson, Timothy. "Metcalfe's Law as a Model for Bitcoin's Value." *Alternative Investment Analyst Review* 7, no. 2 (2018): 9-18.

PlanB. "Modeling Bitcoin Value with Scarcity." *Medium.* 2019. https://medium.com/@100trillionUSD/modeling-bitcoins-value-with-scarcity-91fa0fc03e25.

Popper, Nathaniel. "Warning Signs About Another Giant Bitcoin Exchange." *The New York Times*, November 21, 2017.

Protos. "Tether Papers." Investigation summary, 2021.

Ranger, Matt. "An Anatomy of Bitcoin Price Manipulation." *Single Lunch.* January 2022. https://www.singlelunch.com/2022/01/09/an-anatomy-of-bitcoin-price-manipulation/.

—. "The State of Stablecoin Risks." *Single Lunch.* June 2022. https://www.singlelunch.com/2022/06/14/the-state-of-stablecoin-risks-june-2022/#note-2-2038.

Russo, Camila. *The Infinite Machine.* Harper Business, 2020.

Schär, Fabian, and Aleksander Berentsen. *Bitcoin, Blockchain and Cryptoassets.* The MIT Press, 2020.

Szabo, Nick. "Antiques, time, gold, and bit gold." *Unenumerated.* 2008. https://unenumerated.blogspot.com/2005/10/antiques-time-gold-and-bit-gold.html.

—. "Bit gold." *Unenumerated.* 2008. https://unenumerated.blogspot.com/2005/12/bit-gold.html.

Van Vliet, Benjamin E. "An alternative model of Metcalfe's Law for valuing Bitcoin." *Economics Letters* 165, no. C (2018): 70-72.

Woo, Willy. "Introducing NVT Ratio (Bitcoin's PE Ratio), use it to detect bubbles." 2017. https://woobull.com/introducing-nvt-ratio-bitcoins-pe-ratio-use-it-to-detect-bubbles/.

Zhang, Xing-Zhou, Jing-Jie Liu, and Zhi-Wei Xu. "Tencent and Facebook Data Validate Metcalfe's Law." *Journal of Computer Science and Technology* 30 (2015): 246-251.

Abschließende Erklärung

Ich versichere hiermit, dass ich meine Masterarbeit *„Valuation Models for Crypto Assets: The Case of Cryptocurrencies“* / *„Bewertungsmodelle für Krypto-Assets am Beispiel von Kryptowährungen“* selbstständig und ohne fremde Hilfe angefertigt habe, und dass ich alle von anderen Autoren wörtlich übernommenen Stellen wie auch die sich an die Gedankengänge anderer Autoren eng anlehnenden Ausführungen meiner Arbeit besonders gekennzeichnet und die Quellen zitiert habe.

Münster, den 10. August 2022

____