### NBER WORKING PAPER SERIES

### CRYPTOCURRENCIES AND DECENTRALIZED FINANCE (DEFI)

Igor Makarov Antoinette Schoar

Working Paper 30006 http://www.nber.org/papers/w30006

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 April 2022

Igor Makarov: Houghton Street, London WC2A 2AE, UK. Email: i.makarov@lse.ac.uk. Antoinette Schoar: 62-638, 100 Main Street, Cambridge MA 02138, USA. Email: aschoar@mit.edu. This paper was originally presented at the Brookings Papers on Economic Activity Conference on March 24, 2022. We thank the editor Janice Eberly, and our discussants Gary Gorton and Eswar Prasad for extensive comments that improved the paper. We thank Daniel Ferreira for helpful comments. We also thank Jiageng Liu, Yuou Wu, and Xin Xiong for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Igor Makarov and Antoinette Schoar. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Cryptocurrencies and Decentralized Finance (DeFi) Igor Makarov and Antoinette Schoar NBER Working Paper No. 30006 April 2022 JEL No. G1,G2,G20,G21,G23,G3

### **ABSTRACT**

The paper provides an overview of cryptocurrencies and decentralized finance. The discussion lays out potential benefits and challenges of the new system and presents a comparison to the traditional system of financial intermediation. Our analysis highlights that while the DeFi architecture might have the potential to reduce transaction costs, similar to the traditional financial system, there are several layers where rents can accumulate due to endogenous constraints to competition. We show that the permissionless and pseudonymous design of DeFi generates challenges for enforcing tax compliance, anti-money laundering laws, and preventing financial malfeasance. We highlight ways to regulate the DeFi system which would preserve a majority of benefits of the underlying blockchain architecture but support accountability and regulatory compliance.

Igor Makarov London School of Economics Houghton Street London WC2A 2AE UK i.makarov@lse.ac.uk

Antoinette Schoar MIT Sloan School of Management 100 Main Street, E62-638 Cambridge, MA 02142 and NBER aschoar@mit.edu

The financial system performs a wide array of functions that are important for economic growth and stability, such as allocating resources to their most productive use, moving capital from agents with surpluses to those with deficits, and providing efficient means for moving wealth across time and states, see for example Merton (1995) or Allen et al. (2019). To achieve these goals, the US financial system, and similarly most other countries, have traditionally relied on a set of intermediaries such as banks, brokers, exchanges etc. that are connected by payment systems. These intermediaries serve as centralized nodes that guard the access to the financial system and provide customers with essential services such as record keeping, verification of transactions, settlement, liquidity, and security. This architecture implies that intermediaries perform many of the core functions in the system, and also help with the implementation of regulatory goals such as tax reporting, anti-money-laundering laws or consumer financial protection. As a result, however, these intermediaries can hold significant power, based on their preferential access to customers and data. This centralized position, if not properly harnessed and regulated, can be a source of outsized economic rents and can lead to considerable inefficiencies. It can also lead to inherent fragility and systematic risk if core intermediaries become corrupted or investors lose trust in the system.

The concern about the power and potential corruptibility or fragility of intermediaries, possibly heightened by the experience of the 2008 financial crisis, has contributed to the new "revolution" brought about by the blockchain technology, which is one of the fastest growing financial innovations over the last decade. Its attraction lies in the ability to build decentralized and open access platforms that reduce the reliance on centralized trusted intermediaries and middlemen.

Eliminating unnecessary intermediaries can potentially be a significant benefit of the blockchain architecture. Technological innovations have, of course, long been consequential in improving the efficiency of the financial system or strengthening competition. We can think of innovations like mobile banking or algorithmic lending. What differentiates the blockchain from past technological innovations is that it offers the possibility of a completely different financial architecture, commonly called decentralized finance (DeFi), where record keeping is decentralized, access to the system is anonymous and unrestricted, and any form of intermediation would be built on top of it.<sup>1</sup>

To assess the potential benefits and challenges of the proposed new architecture, it is important to recognize that intermediaries are not merely gate keepers which have no economic value except for rent extraction. Many problems with existing intermediaries

<sup>&</sup>lt;sup>1</sup>DeFi is also distinct from the generic umbrella term of "Fintech". While Fintech innovations also introduce new technologies to financial services, for example Rocket Mortgage which uses on-line origination in mortgage lending, they still rely on a model of centralized intermediaries.

originate from the economic forces that are an inherent part of financial markets and therefore, equally exist in DeFi solutions, but might be relocated to different layers in the new infrastructure as we will discuss. In addition, some of the rents that financial institutions enjoy in the current financial system, are a deliberate regulatory choice: In order to provide institutions with the incentives to abide by regulations, rule-makers allow these institutions to earn some rents to ensure that they have a franchise value.

Advocates of DeFi solutions argue that financial services are ripe to undergo dramatic and disruptive changes. How this evolves, in terms of technology, regulation, and ultimately liquidity and credit to the economy, has important consequences for the US and global economies. There are also strategic and competitive implications across countries. The goal of this paper is to raise some of the issues that arise in a system of Decentralized Finance and propose some solutions, while at the same time providing an introduction to how such a system works and the mechanics behind it.

We start by laying out how the blockchain technology works that underpins virtually all DeFi solutions. We discuss the different ways security is achieved under different protocols, in particular Proof of Work (PoW) and Proof of Stake (PoS), and what economic incentives are built into these solutions to ensure the integrity of the blockchain ledger. Our analysis highlights that the current security protocols have in-built economic incentives for concentration of mining or validator capacity due to inherent fixed costs and benefits of co-insurance for validators. We also show that large PoW networks can have negative externalities on the security of smaller PoW networks that has important implication for the competitiveness of PoW protocols. For PoS platforms, an added complexity arises from the fact that the going concern value of the platform also affects the security of the platform itself and applications that run on it.

Next, we discuss the benefits and limitations of smart contracts. These are selfexecuting pieces of scripting code that can in theory carry out any computation and are the building blocks of many DeFi applications. Since smart contracts are designed not to have recourse to the legal system, they have to be written as complete contracts upfront. We highlight the implication of such a change in the enforcement of contracts on the transaction costs of writing contracts, the opportunity of opting out of current remedial laws, and challenges for consumer financial protection if smart contracts are written outside typical legal protections. Many of these challenges might give rise to a new layer of "trusted" intermediaries, in particular coders who will help people to navigate the DeFi infrastructure that might be too complicated for individual participants. In this context, we explain the role and design of Oracles, which provide access to data from outside the blockchain and allow smart contracts to interact with the real world. Based on these building blocks we then provide an overview of the current crypto landscape and the main DeFi applications such as decentralized crypto exchanges, borrowing and lending markets, and yield farming.

Finally, we compare this new DeFi architecture to traditional financial market solutions and lay out how these two regimes solve some of the most important problems in financial systems, such as data privacy and transparency, extraction of rents, transactions costs, governance issues and systemic risk.<sup>2</sup>

DeFi applications might have the potential to democratize finance by creating a level playing field among providers of financial products and services. But we show that the current design of DeFi applications, which are predominantly built on permissionless and pseudonymous blockchains generates formidable challenges for tax enforcement, aggravates issues of money laundering and other kinds of financial malfeasance, and as a result creates negative externalities on the rest of the economy. Similar to the traditional financial system, there are several natural points where rents can accumulate at different layers in the DeFi architecture due to endogenous constraints to competition caused by network externalities and economies of scale. Also, rent extraction can be driven by frictions at the customer level due to lack of financial sophistication or behavioral biases. In cases where market competition does not work to restrict excessive rents, typically regulations are used to protect the interest of users. But here again the permissionless and pseudonymous design, severely limits the ability of regulators to restrict unscrupulous operators.

The pseudonymous and permissionless structure also has implications for the governance of DeFi apps. Many DeFi apps in their quest to avoid placing trust in any actor or institution have experimented with new organizational forms, so-called decentralized autonomous organization (DAO). The basic idea of DAO is to spread control over decisions among all interested stakeholders by issuing special "governance" tokens that give their holders the power to propose changes to the protocol and vote on them. We discuss the governance challenges that arise in such arrangements, and show that they face the same fundamental governance issues as traditional organizations. As a result, we show that in the majority of crypto projects ownership is concentrated.

Lastly, we discuss the potential of DeFi solutions to contribute to systemic risk and have spillover effects on the rest of the economy. We highlight that DeFi so far has operated under a narrow banking model. This removes many of the problems faced by the fractional reserve system, but also constrains the efficient use of capital. Presently, the main systemic risk comes from the ability of investors to take highly leveraged and

<sup>&</sup>lt;sup>2</sup>Harvey et al. (2021), Schär (2021), Aramonte et al. (2021) also provide detailed discussions of the DeFi eco-system and its applications.

interconnected positions and a potential run on stablecoins. So far, the systemic risk has been limited, but as ties between the regular financial system and DeFi increase the risk can grow.

We conclude by discussing challenges and potential solutions for regulators and market participants in this new infrastructure. A natural place for regulatory oversight in this new ecosystem is at the level of developers and validators, which in turn control the network protocol. Once this level of regulatory compliance is established, many other functions can be built that would address the majority of issues we outlined above. This solution looks similar to a permissioned blockchain, but it preserves most of the desired properties of the blockchain such as observability of transactions, automatic settlement, and execution of the same set of smart contracts.

If regulators give up on the ability to oversee validators, the effectiveness of regulation will be much more limited and will depend on the goodwill and voluntary cooperation of validators and developers of the blockchain. If validators accept transactions from every party, the most regulators could hope for is to separate the network into "regulated" and "unregulated" parts. The latter part could then harbor bad actors and facilitate illegal activities. The opportunities of sidestepping the regulated part will generally increase with the level of crypto-adoption, since people will be able to predominately transact in the unregulated part and avoid triggering regulatory oversight.

## 1. Blockchain technology

A typical financial system can be represented, at an abstract level, as a collection of states and transactions that describe the transition from one state to another. For example, in a payment system a state is a collection of all the accounts in the system together with their balances. Transactions specify how funds move between accounts.

Historically, financial intermediaries have been the key nodes in the financial system that control the accuracy of customer accounts, perform bookkeeping functions and ensure that unauthorized persons do not have access to an account. For a long time, this centralized model of bookkeeping was the only viable option. But recent advances in technology have enabled an alternative architecture of storing and managing information where no single entity has full control over all the states and transactions, or any subset of them. Instead, multiple parties (validators) hold their own copies of states and jointly decide which transactions are admissible. This architecture became known as Distributed Ledger Technology (DLT). A blockchain is a form of DLT, in which all transactions are recorded and organized in blocks that are linked together using cryptography. Bitcoin was the first and remains the most famous application of the blockchain technology.

One of the main advantages of DLT is the elimination of a central point of failure. Since multiple copies of records exist, the corruption of a single node or a single copy has no effect on the security of the blockchain. In fact, the blockchain protocol allows for multiple points of failure or corruption as long as the majority of validators are not corrupted. In particular, it allows validators to be parties that do not trust one another or are even adversaries.

Blockchains are usually divided into permissioned and permissionless ledgers depending on the set of entities that are allowed to be validators. In a permissioned blockchain, a set of validators is approved by a coordinating body, which can be a private firm or a consortium of institutions. In contrast, a permissionless blockchain does not impose ex ante constraints on the number or identity of validators. In addition, blockchains are sometimes categorized as private or public ledgers. In a public blockchain, everyone has full access to the information stored on the blockchain. In contrast, only authorized parties can observe transactions in private blockchains. Typically, permissioned blockchains are private, and permissionless blockchains are public.

Permissioned blockchains still require trust in the coordinating body that approves validators, which is viewed by many crypto enthusiasts as a fundamental flaw. In contrast, permissionless blockchains do not rely on trust in any individual validator, forming what famously has been called a "trustless" trust architecture. The "trustless" trust, however, comes at a high cost. Since anyone can become a validator in a permissionless blockchain, the system is potentially vulnerable to a Sybil attack where an adversary subverts the system by creating a large number of pseudonymous validators and uses them to gain disproportionately large influence over the consensus protocol.

Two main approaches have been proposed for permissionless protocols to be resilient to a Sybil attack, (1) Proof-of-Work (PoW) and (2) Proof-of-Stake (PoS). The main idea behind both approaches for validating transactions is to provide validators with a reward for their services and to make it costly for an adversary to attain a majority stake and subvert the system. The reward is meant to provide validators with financial incentives to work honestly. The reward usually comes in two forms: transaction fees and a pre-specified amount also known as a block reward. The block reward is typically denominated in the platform's native currency and is financed through issuance of new coins, thus serving as a dilution tax on all users.

The decentralization of the ledger also has implications for the scalability of the network. Intuitively, as the ledger becomes more decentralized more copies need to be distributed and more resources need to be spent to achieve the protocol consensus and make the blockchain secure. This trade-off between decentralization, security, and scalability was famously formulated by Vitalik Buterin, a co-founder of Ethereum, in the early days of Ethereum and became known as the scalability trilemma (or sometimes as the blockchain trilemma). The trilemma has attracted a lot of attention and a large number of new blockchain solutions have being introduced to achieve simultaneously the three goals.<sup>3</sup>

In the following, we leave aside the technical issues such as scalability. We also refrain from a game-theoretic analysis of security of different protocols. For an example of such analysis see Biais et al. (2019) and Halaburda et al. (2021). Instead we focus on the embedded economic mechanisms and incentives that are at the heart of the different protocol security approaches. Since most DeFi applications are currently built on permissionless blockchains we will focus predominantly on these blockchains. We show that both PoW and PoS favor validator concentration, since there are strong implicit incentives for validators to pool their capacity and co-insure their risk of winning a block reward. We also discuss the resilience of PoW and PoS to an attack and show that large existing networks have negative externality on small networks. These properties have important implications for competition in the crypto space which we discuss in Section 4.

## 1.1. PoW protocols

In a PoW protocol such as Bitcoin, validators (also known as miners) compete for the right to verify transactions and obtain their reward by solving a computationally intensive problem. For a successful attack on a blockchain an attacker needs to control a large fraction of the total network power, typically 51%, which resulted in the nickname "51% attack". Once an attacker controls the majority of mining power they can alter transactions in the system, for example, they can spend the same cryptocurrency multiple times (known as a double-spending attack).

The likelihood of an attack in a PoW protocol, therefore, depends on the prospects that a malevolent party amasses enough computing power. Notice that miners should at least break even in the long-run to be willing to invest in mining. Therefore, the expected rewards collected for mining a block should cover the cost of its mining. This implies that there are no economic disincentives of amassing 51% and the constraint is on the feasibility of amassing 51% of hashing power, Budish (2018).<sup>4</sup>

 $<sup>^3{\</sup>rm These}$  include sharding, side chains, and lightning networks. There are also non-blockchain solutions, e.g., has hgraph technologies.

<sup>&</sup>lt;sup>4</sup>Hashing power or hash rate is the primary measure of a miner's performance in a PoW protocols. See https://en.bitcoinwiki.org/wiki/Hashrate for more details.

Of course, any successful attack on a blockchain reduces trust in this blockchain, and therefore, its economic value. If miners have to incur large fixed costs to set up their operations, then by attacking the blockchain they will forfeit some of the future profits and might not be able to recover their initial investments. This reduces the benefits of the attack and therefore, can make it unprofitable.

Therefore, the lower the fixed costs, the less costly is a 51% attack. As a result, any factors that reduce fixed costs have negative effects on the security of the network. In particular, large PoW networks like Bitcoin or Ethereum have negative externalities on the security of smaller PoW networks.

The large appreciation of Bitcoin and Ethereum led to significant investments in mining capacity. Smaller networks like Litecoin or Bitcoin Gold usually attract only a small fraction of the mining capacity of these larger coins, since their rewards also are much lower. This creates a possibility that a miner with a large hashing capacity can divert a fraction of it to attack a smaller coin, if they chose to.

Furthermore, the emergence of market places like NiceHash, where mining hash power can be rented for a specific time period, has made it possible for people to speculate on mining profitability without owning the physical hardware themselves and amass hashing power for a possible attack. The amount of available hashing power in these market places is only a small fraction of the capacity used in large networks such as Bitcoin and Ethereum, which usually operate close to full capacity. But the available capacity on NiceHash, often is significantly larger than the total mining capacity employed in smaller networks.<sup>5</sup> These renting opportunities have significantly reduced the cost of a 51% attack on smaller networks, and in fact have led to many such attacks on smaller cryptocurrencies such as Bitcoin SV (BSV), Bitcoin Gold (BTG), and Ethereum Classic (ETC), see Table A.1 in the Appendix.

The negative externalities of large PoW networks on smaller networks has important implications for the competitiveness of PoW protocols. It suggests that once one or a few major PoW blockchains are in existence, new entrants might find it difficult to compete. While the new protocol has not reached a critical mass yet, it has a heightened likelihood of being subject to an attack. This makes it less less secure and therefore, might reinforce the dominant position of the first movers. One defense against the negative externalities of hashing capacity in larger blockchains would be to make mining equipment very platform specific, so that slack in a larger system does not affect the new entrant. However, platform specific mining hardware can increase entry cost for miners to the new platform and therefore, can have negative effect on its

<sup>&</sup>lt;sup>5</sup>See, for example, the website Crypto51 https://www.crypto51.app/ that measures the cost to 51% attack Bitcoin and other major proof of work cryptocurrencies.

growth and security.

While there have not been any successful 51% attacks on Bitcoin or Ethereum, it does not mean that they are completely safe from it. First, as we mentioned above, these networks have benefited so far from large price appreciation that have made miners operate at nearly full capacity. If at some point there is a substantial price decline it is likely that an increasing number of miners will find it unprofitable to continue their mining operations. This can lead to an increase in spare mining renting capacity, and therefore, might increase the probability of an attack.

Second, the original design of Satoshi Nakamoto, the inventor of Bitcoin, envisioned a world where mining would be fully decentralized and not depend on a few large players. In this world, miners would find it difficult to collude and failure of any one miner would have no consequence for the security of the network.

This original idea, however, clashes with the economics of mining in PoW protocols. By design, the probability of winning the race and obtaining the block reward is proportional to the computing power spent on mining. This gives strong incentives for miners to pool their computing power and co-insure each other. As a result, mining in most PoW blockchains is dominated by large mining pools, Cong et al. (2020), Ferreira et al. (2019).

The concentration of mining pools has attracted a lot of public attention and concern, since high concentration facilitates collusion among miners and with it the danger of an attack. Even if miners themselves do not misbehave, high concentration increases the risk that a malevolent party, either a private or a state actor, could hijack them and gain control over the network.

Some observers downplayed the risk of the attack coming from pool concentration arguing that even though pools can have substantial influence over the cryptocurrency protocol, they do not necessarily control their miners. Therefore, if any pool is noticed to engage in rogue behavior its miners can leave it and join other pools.

The power that a pool operator has vis a vis individual miners depends on the ease with which miners can shift capacity across pools, which in turn depends on the underlying size distribution of the miners. Makarov and Schoar (2021) document that miner concentration in the Bitcoin protocol is high, even at the level of individual miners. They show that at times fewer than 50 miners control 50% of mining capacity. One explanation for this concentration in mining power seems to lie with the high fixed costs of setting up a large mining farm that result in increasing returns to scale.

The paper also shows that the concentration of mining capacity is counter cyclical and varies with the Bitcoin price. It decreases following sharp increases in the Bitcoin price and increases in periods when the price drops. Thus, the risk of a 51% attack increases when the Bitcoin price drops and makes the system more fragile.

### **1.2.** PoS protocols

While the costs of an attack and the resilience of a PoW network increase with the size of the network so does the cost of verification. According to the Cambridge Bitcoin Electricity Consumption Index, the annual electricity consumption of the Bitcoin network in 2021 reached 130TWh, which exceeds the annual consumption of such countries as Norway or the Ukraine. Because miners have to be compensated for their costs large electricity consumption translates into high transaction fees. Figure 1 shows the average transaction fees in the two largest PoW protocols, Bitcoin and Ethereum. As the Bitcoin and Ethereum prices have significantly increased over time and so have the fees.

## [Fig. 1 About Here]

The serious concerns about the sustainability and energy consumption of PoW protocols have favored the emergence of PoS blockchains. PoS protocols consume significantly fewer resources than PoW protocols. Platt et al. (2021) estimate energy consumption of major PoS protocols and show that their energy consumption per transaction is comparable to that in the Visa network. Recognizing the drawbacks of PoW protocols, after 2017 there was a significant acceleration in the development of PoS blockchains Also, Ethereum instituted a shift to a PoS protocol, Ethereum 2.0 (Eth2), to be completed in 2022.

In a PoS protocol, instead of solving a difficult mathematical problem a validator stakes its coins, which can be forfeited if the validator fails to verify transactions in a timely manner or its actions are determined to be malicious. In most PoS protocols, participants who stake more coins are more likely to be chosen to verify transactions (or have more rights to vote for a validator in delegated PoS networks). Thus, PoS protocols are built on the idea that a party that has a large stake in the given network would not want to undermine this network since the gains from an attack would not compensate for the loss of value that comes from penalties and the drop in the network's valuation.

The above argument relies on the idea that a validator which owns a large stake in the platform also has an interest in its continuation value and thus should be disincentivized from endangering it. This logic makes sense, if the attack in question is, for example, a double-spending attack, since the gains in that case are a small fraction of the total value of the network. However, the gains from an attack might not be restricted to simple gains from double-spending. First, if the network is part of a competitive environment, competing networks might realize substantial gains from undermining a new entrant. Similarly to what we described in PoW blockchains, here as well undermining fledgling rivals can be particularly profitable if it reduces future competition.

Second, many PoS blockchains are smart contract platforms that position themselves as a base layer that provides security for other applications or even other blockchains that are built on it. In this case, there is tension between the value of the base layer blockchain and its applications. If the value of the base layer is below the value of an application, an attacker who wants to undermine the application might find it profitable to attack the base layer. To prevent such an attack the value of the blockchain at the base layer should be substantially greater than the value of its applications. Since the value of the base layer comes primarily from transaction fees (and seignorage), the possibility of an attack on the base layer puts a lower bound on the required size of the fees that have to accrue to the blockchain at the base layer. High fees, however, hurt the value of applications built on the platform, and thus the platform's value.

We showed in Section 1.1 that mining in PoW blockchains is dominated by pools because they allow miners to co-insure each other. A similar force is at play in PoS blockchains. Since the probability of being chosen and collecting the reward depends on the amount of coins a validator is staking, investors have incentives to pool their stakes together and co-insure each other.

Table 1 documents concentration of validators for the largest PoS protocols as of January 2022. The data shows significant concentration for the vast majority of the PoS blockchains. The top 10 validators hold typically more than 25% of the capacity, while the top 50 validator are above 50%.

In addition, since the technology used across different PoS protocols shares many similarities the same validators typically work on multiple blockchains. Table 2 shows the top 15 validators together with their combined stakes in the top 10 largest PoS protocols. The top 10, 50, 100 validators account for 14%, 32%, and 41% of stakes across the 10 largest PoS blockchains, respectively.

The concentration of PoS validators as of now is lower than in the PoW protocols, but it is not fully dispersed either. It is of interest that a few validators are starting to emerge as dominant players across different blockchains.

## 2. Smart contracts

Smart contracts have become another fundamental layer of the new DeFi architecture. To go beyond simple interactions such as the transfer of coins or assets on the blockchain, many newer protocols starting from Ethereum provide the opportunity to embed pieces of scripting code that can in theory carry out any computation. These pieces of code became known as "smart contracts". The name and the concept are credited to the cryptographer Nick Szabo who defined smart contracts in 1994 as "a set of promises, specified in digital form, including protocols within which the parties perform on these promises". The modern implementation of this idea arrived with the creation of Ethereum, which is designed to execute smart contracts and makes it convenient for developers to build applications on top of the blockchain.<sup>6</sup>

By itself using software code to represent and execute contractual agreements is not new. For example, when trading via an online brokerage platform, each time a customer sets up a limit order that automatically buys certain stocks when prices match a predefined level, the contract is executed by a software program. Financial markets or e-commerce are dominated by these types of arrangements since they allow a large volume of transactions to be executed quickly and efficiently. But even if the program automatically executes a set of tasks, in traditional electronic contracts, the parties to the contract still have recourse to the legal system if there is a dispute. For example, if a limit order is executed based on wrong information used by the online brokerage platform, the client can seek restitution from the brokerage through the courts.

The critical differences, from an economic perspective, between traditional electronic arrangements and smart contracts that are executed on a permissionless blockchain, arise from how the contracts are executed and enforced.<sup>7</sup> We show since smart contracts are self-executing once they have been embedded in the blockchain, they require contracting parties to complete contracts as much as possible ex ante, since they cannot rely on the ex post remedial protections through the legal system. We discuss the implications of this switch for the transactions costs of writing contracts, the ability of contracting parties to opt out of the current legal protections, and the constraints to consumer financial protections. The need to import up-to-date information from the outside (off-chain) world into the blockchain also led to the development of a new set of entities, so-called oracles. We lay out their role for the functioning of smart contracts

<sup>&</sup>lt;sup>6</sup>See Ethereum White Paper, https://ethereum.org/en/whitepaper/.

<sup>&</sup>lt;sup>7</sup>Smart contracts can also be implemented on permissioned blockchains. In this paper, we focus on smart contracts run on permissionless and public blockchain protocols, since their major applications have been hosted on such blockchains.

and potential vulnerabilities that are introduced through oracles. Finally, we argue that this new architecture might require contracting parties to rely on a new set of trusted intermediaries, such as the developers of the smart contract platform or coders that help to write the computer programs that will be executed on the blockchain.

## 2.1. Execution and enforcement

The execution of a smart contract on a permissionless blockchain fundamentally changes the process of enforcement, see Werbach and Cornell (2017), Werbach (2018). First, once a program has been executed, the distributed nature of the contract verification makes it impossible to unilaterally stop or reverse its execution, unless certain conditions for stopping the smart contract were included in the program ex ante. Second, even if one party wanted to sue a counterparty, there might not be any party that can be held accountable because of the anonymity of the transactions. Practically speaking there might be no one who can be served with a legal notice.

These changes are important for the application of contract law, since it is fundamentally a remedial institution that operates on an ex-post basis. First, contract law aims to rectify situations ex post, where one party has wronged another party by breaching the terms of the contract or not delivering on a promised action.

Second, the law incorporates a variety of doctrines, which allow one or multiple parties to annul the contract ex post. These exemptions are meant to protect contracting parties against unwittingly (or deliberately) taking advantage of each other or of an unforeseen situation. These are issues such as unconscionability, mutual mistake, illegality, capacity, consideration, fraud, or duress. The role of judges and the legal system then is to oversee and enforce the intended application of the law in these cases. In other words, the role of the legal system is to complete contracts that were either deliberately or unintentionally left incomplete ex ante, see Wright and De Filippi (2015).

Of course, contracts are written in the shadow of the law. The expectations that contracting parties have about how laws will get enforced, affect how contracts are written in the first place, and which parts can be left unspecified. Since smart contracts do not allow for recourse to the legal system, they have to be written as complete contracts upfront. Or at a minimum the contracting parties have to specify exactly which states of the world they are willing to leave unspecified. Since the smart contract cannot be unilaterally stopped and renegotiated, if a state of the world is not ex ante specified, the program will execute as if this state never existed.

This example highlights that a contract breach in the traditional sense is not possi-

ble on the blockchain. Once the parameters encoded in the smart contracts are realized, the code will execute the transaction. This significantly reduces the chance of one party to a contract reneging on it after the fact, say because they changed their mind or they were not serious about the transaction in the first place. But the automatic execution of smart contracts also eliminates the opportunity for "efficient breach". Take the situation of a mutual mistake, where a buyer and seller agreed to the purchase of an asset at a specific price. But just before the seller is supposed to deliver the asset, she realizes that the asset is worth much more than either side had realized. Here, in a traditional contracting situation, the seller could engage in "efficient breach" and not deliver the asset until both side had a chance to renegotiate the terms of the deal. However, with a smart contract the transfer will be executed, since the parties by definition did not plan for the mutual mistake upfront. And a similar logic holds for many of the other protections that traditional contract law awards. This shifts the status quo of which parties will be in the role of plaintiff and defendant.<sup>8</sup>

## 2.2. Smart contract trade-offs

#### 2.2.1. Transaction costs of contracting

As the discussion above highlights smart contracts must be written in precise, fully defined computer code since they cannot be modified once executed. Many proponents of smart contracts have suggested that this reduces their cost, since there is no scope for ex post renegotiation. But these cost savings might be offset by the higher upfront costs of negotiating and specifying the precise terms of an agreement in all possible states of the world. These up-front costs will become especially high when there is large uncertainty about the future states of the world or if these states are hard to imagine and to define ex ante.

To mitigate these issues, traditional contract law systems provide a series of mandatory and default positions that allocate risk when matters are left unspecified. In the case of smart contracts this default to the legal system is not possible. So, the costs must be borne by the individuals engaging in the contract. In the case of contracts, which are very simple and standardizable, some templates of code will most likely be developed that anyone can use to embed in a smart contract. This can reduce the upfront cost in cases where many people have very similar contract issues and the future states and outcomes over which the contract needs to be defined are also very standard

<sup>&</sup>lt;sup>8</sup>Parties to a smart contract could try putting in protections against mutual mistakes by writing into the contract arbitration of third party experts, but this would require trust in experts and therefore, would go against the main idea of smart contracts.

and simple to understand. However, as soon as there is more variation in possible contract templates that can be considered in a contracting situation, the mental cost of comparing and understanding the different options might become quite high. And of course, the costs are even higher if the situation is unique, and a lot of value is at stake. Here parties cannot choose from existing templates but have a strong incentive to not inadvertently miss or miscode a possible state of the world. This means they do have to bear the upfront costs of trying to write as complete a contract as possible.

#### 2.2.2. Smart contracts as a commitment device

Even people who trust the legal system might in some situations want to avoid ex post litigation risk to bring down ex ante cost, for example, reducing the possibility of opportunistic behavior or efficient contract breach ex post. Take a situation where both parties to a contract are well informed about the functioning of a certain financial product, say a mortgage, and thus ideally the lender would not need to spend time developing education material to inform the borrower about what happens in case of default. However, if the borrower has the right to sue ex post if she was not informed that the lender can seize the property, the lender will be forced to develop training material to prove that the borrower has been informed. An informed borrower and lender might be better off if they could shut off the opportunity for the borrower to sue in case of default. It would eliminate the lender's need to invest in expensive training material which is wasteful in this case. But since the borrower cannot abdicate their right to sue, both parties must bear the cost of the upfront training.

These issues apply in situations where both parties to a contract are sure that they do not value any ex-post protection through contract laws. This requires that both sides must be well informed about the logic of the contract, all the possible ex post outcomes and do not fear the possibility of being taken advantage of. In financial markets this is an important concern since many contracts involve investments in complex and risky products, say for example trading in derivatives. If customers could sue each time a bad state of the world realizes and claim that they were misled about the product, intermediaries would not be able to sell any risky securities. In the US the law has addressed these issues by granting certain exemptions to high-net-worth individuals or people who can demonstrate their knowledge in those products. But it does not provide sweeping exemptions from the ex post protections of contract law, since in many situations consumers might not even be aware of their own lack of knowledge relative to an informed market participant.

#### 2.2.3. Smart contracts and consumer financial protection

A large literature in finance has shown that many participants in financial contracts, especially retail investors, lack financial literacy and are not well prepared to understand financial markets, see for example Lusardi and Mitchell (2007). Although parties are generally free to enter into agreements, subject to certain limitations and exceptions, the law protects parties in certain situations by determining whether they had the capacity to enter into a legally binding agreement. For example, contracts may be voidable if made by a minor or persons who are mentally ill or intoxicated at the time of contracting. By not allowing mandatory ex post protections through the legal system, smart contracts do not provide sufficient safeguards for financially less informed or more fragile customers. Since smart contracts typically have limited means to test for a person's financial sophistication or mental capacity, the enforcement of these contracts could lead to undesirable outcomes, if there is no provision to reverse the outcome as in traditional contract law.

If financially less sophisticated consumers are aware of their lack of knowledge and understand that there is a risk that in such an environment they are disadvantaged, the most plausible result would be to opt out of this contracting environment. However, if smart contracts became the predominant form of contracting it would severely affect market participation of less sophisticated consumers. Or alternatively these customers would have to find trusted intermediaries to act on their behalf. So we are back to the original problem of how to ensure good performance of intermediaries. But given the pseudonymity of the blockchain environment it would be more difficult to build trust. Furthermore, a large literature in behavioral finance has shown that many financially unsophisticated consumers are not aware of their lack of information or are overoptimistic about their ability to participate in financial market. As a result they might unknowingly sign contracts that are against their own interests, see for example Laibson et al. (2007) or Campbell (2016).

To curtail the most egregious abuses in the traditional system, the US has a set of consumer financial protection regulations in place such as the Consumer Financial Protection Act, Fair Debt Collection Practices Act, or Truth in Lending Act. These aim to reduce the asymmetry in knowledge and information between financial institutions and customers to provide better outcomes for consumers. The example above suggests that these types of regulations will be difficult to implement on a permissionless blockchain.

#### 2.2.4. Are smart contracts really "trustless"?

An often-highlighted promise of smart contracts is that they may reduce the need for trust between contracting parties, or trust in the legal system. Legal enforcement of contracts can be cumbersome and prone to error. In some societies the legal system itself can even be corrupt and biased. If people do not trust the legal system, they might prefer a decentralized execution that is not subject to ex post discretion. But it is not clear whether trust can be removed altogether from the process of smart contracting or whether it simply requires a shift of trust to other intermediaries and systems.

In a narrow set of circumstances, smart contract can automatically enforce transactions if all parts of the transaction are on-chain. For example, a contract that exchanges one token for another on the same blockchain does not rely on enforcement or adjudication outside the blockchain. Here the level of trust is as high as the trust in the blockchain itself. Some level of trust is still required. For example, parties need to trust the developers that oversee a network's protocol that they have not embedded errors in the coding of the platform. Or that the consensus protocol is well enough designed that it is not prone to any attacks.

However, the vast majority of important financial interactions rely on assets, actions, or information that exist outside the blockchain. For example, one of the most important financial contracts a typical household in the US holds, is a mortgage against their house. While one could imagine a smart contract that uses the home as collateral, the transfer of the house cannot be fully automated on the blockchain ledger. First the smart contract would have to stipulate how the deed record in the public data base must change, in case of default or non-repayment of the loan. Second, even if we assume that the deed record itself lives on the same blockchain, if the person who currently occupies the house does not move out when the ownership changes, it does need off chain verification and enforcement to change the de facto state that matters, e.g. can you occupy the house you supposedly own.

Getting off-chain data presents a number of challenges. The solution revolves around the use of "oracles" — an off-chain entity that creates a transaction on-chain with the data posted. Oracles define how a smart contract incorporates off-chain information into the execution of a program, which we discuss in detail in the next Section 2.3. The consequence of using oracles is that parties need to trust them.

In addition, given the lack of an ex-post appeals process via the law, a lot is at stake when specifying a smart contract to be as complete as possible upfront. Especially for transactions that are more complicated, the machine-readable code for the smart contract must be complete and follow strict rules of syntax and semantics. In practice, most people are not able to write this type of contract themselves, and therefore must rely on coders, or third-party developers. This can lead to perverse incentive for developers who are more knowledgeable than the principal who hires them to take advantage of them and exploit deliberate vulnerabilities in the code. The fact that the code underlying the contract is stored on the blockchain and publicly accessible alleviates but does not completely eliminate the problem. The pseudonymity of the blockchain makes it difficult to confirm if a developer of a code is also the agent benefiting from any vulnerability. And at least currently, developers are not bound by the same fiduciary standards as financial intermediaries.

#### 2.2.5. Observability

When interacting with a regular server-based web application, the user often cannot observe the details of the application's internal logic. As a result, the user has to trust the application service provider. Smart contracts mitigate this problem and ensure that an application runs as expected, since the code underlying the contract is stored on the blockchain and publicly accessible. However, this type of observability can also have a downside if it leads to strategic behavior. For example, take any rating system in finance such as a personal credit score or a firms bond rating. If the smart contract spells out exactly how the score is calculated, users might optimize against the code so that they land just above the cut off for the best category. This could undermine the usefulness of these types of scores, see for example Berg et al. (2020) for an example of loan officers gaming a scoring threshold.

Another possible problem with the observability of data on the blockchain has been highlighted in Cong and He (2019). Since generating decentralized consensus entails distributing information, it changes the information environment for the market participants. In particular, as Cong and He (2019) argue it can encourage greater collusion between interested parties.

## 2.3. Oracles

While the blockchain tries to remove the reliance on third party enforcement, smart contracts often need to access data from outside the blockchain if they want to interact with the real world. Consider, for a example, a limit order, where a person writes a smart contract to automatically sell a token of Bitcoin when the price hits a certain target level. For this contract to work, the contract needs to access up-to-date Bitcoin prices. If the data is not obtained in an accurate and timely fashion, a smart trader could reap large gains by taking advantage of stale or wrong prices. One solution would have been to allow the smart contract to obtain the price by querying an API of some exchange. The problem with this solution is that almost all blockchains are designed to be deterministic, which means that any state should be reproducible given the history of the network transactions. Determinism is important so that different nodes that execute the contract can come to a consensus. Since querying the internet can in general produce different values (for example, the price depends on the time of the query), allowing the smart contract to query the price would lead to different values across the nodes, thus making the consensus impossible.

A solution to the above problem is to use an off-chain entity that does the query and posts the data on-chain. Once the data are on-chain smart contracts can access and use them. The off-chain entities that query, verify, and authenticate external data sources and then transmit the information to a blockchain in the crypto parlance are called oracles.

There are many types of oracles.<sup>9</sup> The central issue in design of any oracle is trust. Similar to a chain, which is as strong as it is weakest link, a smart contract is as secure as its least secure components. If the data supplied by an oracle are corrupted then so is the output of the smart contract.

The simplest design of an oracle is where an entity queries a single data provider and records the data on the blockchain. For example, it could be a query from a Coinbase web API. This is called a centralized oracle, which is often a fast and efficient solution. However, reliance on one centralized entity and one centralized data source introduces several potential points of failure. First, the entity can be corrupted. For example, the oracle could withhold the data or front run on information it provides. Second, the data can be corrupted in the process of transferring from the data source to the blockchain because of a software bug. Finally, the data source itself needs to be trusted.

In its perpetual quest to minimize trust from relying on third parties, the crypto community therefore has been actively working on new oracle designs. Inspired by the decentralized trust model of permissionless blockchain protocols, decentralized oracles have become one of the fastest growing solutions, with Chainlink currently dominating the space. The main idea behind any decentralized oracle is (a) to source data from a large and heterogeneous set of entities (nodes) to determine the validity and accuracy of the data and (b) by using incentive mechanisms and skin-in-the-game to keep the entities honest.

Similar to PoS protocols, every participating node that delivers data has to stake

 $<sup>^{9}\</sup>mathrm{See},$  for example, Beniiche (2020) and Caldarelli and Ellul (2021) for a survey of different oracle types.

a deposit, typically in the native token of the network. If the node provides accurate data, it earns a reward. If it misbehaves, the node can lose a percentage of its stake, and in some cases, access to future participation in the oracle network and as a result, all future revenue from the protocol.

The fundamental challenge then is to determine what the truth is. In a blockchain, the correctness of transactions is a property of internal consistency (no double spending). There can be multiple conflicting versions of the blockchain (forks), but there is always one that is correct, and the goal of validators is to agree on which one. In an oracle network, the situation is more complicated. Depending on the nature of the data in an oracle network, there might not be the "true" report, but only its noisy realizations. Therefore, a typical solution to determine the consensus report is to rely on the wisdom of the crowd and use some form of aggregation across reports, e.g., take the median or mean value.

This reliance on a diversified set of data providers, however, exposes the process to the possibility of an adversarial attack, where an adversary bribes the existing nodes or sets up nodes to produce a corrupt report. Equally problematic, could be collusion among oracle nodes. If the gains from collusion become very high, the oracle nodes might not care to lose their current stakes or even all future stakes. As a consequence, the oracle's economic rent should be high enough to ensure that its members are to remain honest.

The research on decentralized oracles is in fledgling state.<sup>10</sup> There are many open questions. For example, holding the size of oracles network fixed what design is the most resilient to the bribery attack? Is it optimal to restrict the size of the network or allow a free entry of nodes? Holding economic rent of an oracle fixed, what is the maximum stake that can be written on the oracle's output?

## 3. The current cryptocurrency landscape

According to CoinGecko, there are currently over 10,000 crypto tokens with an aggregate market cap of more than \$2 Trillion. Several classifications have been proposed for crypto tokens.<sup>11</sup> We have found useful to parse the universe of crypto tokens into the following large categories depicted in Figure 2.

[Fig. 2 About Here]

 $<sup>^{10}</sup>$ See the Chainlink white paper https://chain.link/whitepaper and references there in.

<sup>&</sup>lt;sup>11</sup>See, for example, Cong and Xiao (2021), Prasad (2021). The recent CAP report describes how cryptocurrencies fit in the current regulatory landscape, https://www.americanprogress.org/article/congress-must-not-provide-statutory-carveouts-for-crypto-assets/.

## 3.1. Stablecoins

To start with, we can separate crypto tokens into stablecoins and non-stablecoins. Stablecoins are designed to maintain a peg to fiat-currencies and therefore, act as a safe asset that is not subject to the same volatility as many cryptocurrencies. The absence of central bank digital currency (CBDC) and the growth of DeFi applications based on smart contracts created a strong demand for private stablecoins that are native to cryptocurrency protocols. If at the beginning of 2021 the market value of all stablecoins was \$30 billion then by February 2022 it has reached \$180 billion. As a point of comparison, the total amount of British Pound banknotes in circulation is equal to about £80 billion.<sup>12</sup>

The existing stablecoins can be divided into stablecoins backed by traditional liquid and safe assets, for example, US dollars and treasury bills, and algorithmic stablecoins backed by other cryptocurrencies. Figure 3, the left panel shows the relative share of stablecoins backed by traditional and crypto assets, with the former being the vast majority.

#### [Fig. 3 About Here]

The middle and right panel show the largest stablecoins within each category. The stablecoins backed by traditional assets are dominated by just three coins Tether, USD Coin, and Binance USD. To guarantee the peg, the stablecoins backed by traditional assets should be backed one to one by cash or cashlike assets such as US treasuries. Many stablecoin providers had made claims that their tokens were 100% backed by liquid assets, only later to reveal that it was not the case. The famous examples include the two most popular stablecoins, Tether and USD Coin.<sup>13</sup> In both cases, some part of collateral was held in securities subject to default risk. In October 2021, Tether was fined \$41 million by the Commodity Futures Trading Commission for making misleading claims about being backed one to one by the US dollar.

Along with the stablecoins backed by traditional assets there have also been growing acceptance of algorithmic stablecoins. The combined value of algorithmic stablecoins exceeds \$25 billion, with the largest coins being Dai and Terra USD. The rising popularity of algorithmic stablecoins can again be traced to the desire of crypto community not to rely on centralized parties. Since fiat currencies are issued by governments, the

<sup>&</sup>lt;sup>12</sup>See https://www.bankofengland.co.uk/statistics/banknote.

<sup>&</sup>lt;sup>13</sup>See, for example, https://www.coindesk.com/markets/2021/05/13/tethers-firstreserve-breakdown-shows-token-49-backed-by-unspecified-commercial-paper/ and https://www.coindesk.com/business/2021/07/20/circle-reveals-assets-backing-usdcstablecoin/.

stablecoins backed by traditional assets depend on the trust in government. To break from the need to trust the government, algorithmic stablecoins, or as they are often called programmable money, use other cryptocurrencies as a collateral or/and sophisticated algorithms to regulate the stablecoin supply so that to maintain the peg.<sup>14</sup>

There are now increasing calls for an urgent regulation of the stablecoins. The main concern is that lack of transparency in reporting of the reserves and inadequate collateral can make stablecoins to be prone to a run. We get back to these issues in Section 4.5.

## 3.2. Non-stablecoins

Non-stablecoins constitute a large and diverse group. Their value depends on the current investor sentiment and fluctuates widely over time. First, we can isolate coins that have no other function than being a cryptocurrency, either used for transaction purposes or as a store value. This group includes the first generation of cryptocurrencies such as Bitcoin, Litecoin, etc. By construction, these are the cryptocurrencies that are built on non-smart contract platforms. The majority of these cryptocurrencies are based on PoW blockchains. Early on, crypto enthusiasts hoped that these cryptocurrencies could replace government-sponsored currencies as a transaction medium. However, it quickly became clear that this was infeasible because verifying transactions on public PoW ledgers is slow and highly energy-inefficient. Since then, a new narrative for the benefits of these coins emerged, with them positioned as the new "gold" — a digital store of value. Figure A.1 in the Appendix shows that Bitcoin dominates this group with a market share of more than 90%, followed by Dogecoin. Dogecoin was created in 2013 by two software engineers Billy Markus and Jackson Palmer as a parody of a cryptocurrency that was meant to be worthless. It sharply increased in value and became the first meme coin in 2021 following public support by Elon Musk.

#### 3.2.1. Smart-contract platforms

Another large group are tokens issued by smart contract platforms such as Ethereum, Binance Smart Chain, Solana, Cardano, etc. In many ways, these tokens are similar to the tokens in the first group. In particular, they can also be used to pay for transactions on the platform and are a claim on the platform's economic value. The reason we separate them from the first group is that cryptocurrencies in the first group offer no

<sup>&</sup>lt;sup>14</sup>See, for example, Dai and Terra White Papers for further details: https://assets.website-files.com/611153e7af981472d8da199c/618b02d13e938ae1f8ad1e45\_Terra\_White\_paper.pdf, https://makerdao.com/en/whitepaper.

intrinsic economic value other than the potential for capital appreciation. Therefore, it is unclear what aggregate risk, other than inflation, they are supposed to be tied to.

In contrast, the value of a smart contract platform depends on the scope and the number of applications run on the platform since they affect the number of transactions and the amount of transaction fees, which in turn, influence the price of the platform token.<sup>15</sup> Figure A.2 and Figure A.3 in the Appendix show the development of smart contract platforms. The left panel of Figure A.2 shows the evolution of the market value of different platforms. The right panel shows platforms' market share. Figure A.3 shows the growth of the total value locked (TVL) on the platforms. TVL is the overall value of crypto assets deposited in applications run on the platform. It has emerged as a main metric for gauging interest in a particular platform or sector of the crypto industry.

Figure A.2 and Figure A.3 show that smart-contract platforms grew exceptionally fast in 2021. If at the start of 2021, the total market value of smart-contract platforms was around \$144 billion, at the end of January 2022 it stood at \$683 billion, almost reaching \$1 trillion in November 2021. Similarly, the combined TVL across all platforms was 18 billion in the beginning of 2021 and grew to about 177 billion by February 2022.

Figure A.2 and Figure A.3 also show that Ethereum still dominates this space. The market share of Ethereum has been relatively stable at about 50%. The high fees on Ethereum platform, however, have led to the growth of other smart platforms and to an increase in share of applications deployed on them. If in the beginning of 2021, Ethereum completely dominated the space, by the end of 2021, its share declined to 66%.

#### **3.2.2.** DeFi applications

Smart contracts layered on a permissionless blockchain protocol have given rise to the emergence of what is called decentralized finance (DeFi) — a suite of financial applications meant to replicate many of the elements of traditional financial system without relying on centralized intermediaries.

Figure 4 shows the five largest DeFi sectors. The main applications so far have been centered on trading platforms, borrowing and lending marketplaces, oracles, yield farming, and insurance. The left panel shows the evolution of the market value of the sectors; the right panel — the total value locked in each sector.

<sup>&</sup>lt;sup>15</sup>This division in the two groups is a simplification since even the Bitcoin blockchain can host other protocols, e.g., Omni layer, or help secure other platforms, e.g., Rootstock, DeFiChain. However, presently the scope of these applications compared to those built on smart contract platforms is limited.

#### [Fig. 4 About Here]

**Decentralized crypto exchanges.** Decentralized exchanges (DEXs) have attracted a lot of attention and have become the fastest growing sector of the DeFi universe. One of the main advantages of decentralized exchanges over centralized exchanges is the ability for users to keep control of their private keys. When market participants deposit their crypto tokens with a centralized exchange they forfeit their ownership to the exchange. This makes them exposed to exchange risk — if the exchange is hacked and its funds are stolen investors can experience significant losses. More generally, trading on a centralized exchange requires participants to trust in the exchange that goes against the maxim of decentralized finance. Trading on DEXs is governed by smart contracts and eliminates counterparty risk for the investors. The settlement of transactions is instantaneous, after they are confirmed and included on the blockchain.

The majority of DEXs use an automated market maker (AMM) protocols, which allow a direct exchange of two crypto tokens, say X and Y. The main object in an AMM protocol is a new market structure called liquidity pool. A liquidity pool consists of two pools: one of X tokens and one of Y tokens. The ratio of tokens in each pool defines the current exchange rate between the two tokens.

A liquidity pool supports two main operations: liquidity provision and a swap between the two tokens. Anyone who owns the two tokens can choose to be a liquidity provider by depositing tokens X and Y to the respective pools in the proportion equal to the current ratio. In return, the liquidity provider receives a claim on the share of the two pools' tokens (the so-called LP tokens).

A swap order allows one to exchange one token for the other. The exchange rate depends on a particular implementation of the AMM protocol and is determined by some deterministic rule called the bonding curve. For example, in the constant product AMM used by a popular DEX Uniswap 2 if the initial amounts of X and Y tokens in the liquidity pool are x and y, and someone wants to exchange  $\Delta x$  of X tokens for Y tokens the exchange rate is determined according to the following rule:

$$(x + \Delta x) \cdot (y + \Delta y) = x \cdot y \quad \Leftrightarrow \quad \frac{\Delta y}{\Delta x} = -\frac{y}{x + \Delta x}.$$
 (1)

Swapping X for Y increases the relative share of X tokens in the liquidity pool and therefore, lowers its price relative to the price of Y tokens. Whenever the equilibrium price of the two tokens deviates from the current ratio in the two pools, one can profit from it by executing a swap order until the ratio reaches the equilibrium price. To compensate liquidity providers for providing liquidity everyone who executes a swap order pays a transaction fee that goes to the liquidity pool. This is similar to limitorder book exchanges, where liquidity takers executing a market order usually pay to liquidity providers who supply limit orders.<sup>16</sup>

The DEX's smart contract usually allows to trade any pair of tokens supported by the underlying blockchain. For example, Uniswap 2, realized on the Ethereum blockchain, allows to trade any pair of ERC 20 tokens. If no liquidity pool exists for a particular pair of tokens, it can be freely created. The viability of the pool then depends on the ability of the pool to attract liquidity providers and traders. The liquidity is usually concentrated in a few pairs. Figure A.4 in the Appendix shows how DEX trading volume compares against CEX. While the volume of DEX has experienced fast growth, it still constitutes only a fraction of the CEX volume.

Similar to centralized exchanges, a few DEX dominate the space. Figure A.5 in the Appendix, left panel shows the market share of the top 10 centralized exchanges, the right panel shows the top 10 decentralized exchanges. The majority of CEX volume is concentrated on offshore exchanges such as Binance, Huobi, OkEx, and FTX, which are subject to little or no regulatory oversight. Similarly, Uniswap, PancakeSwap and SushiSwap account for about 70% volume among decentralized exchanges.

**Borrowing and lending.** Lending protocols have been another fast growing sector of the DeFi. Similar to DEXs, lending and borrowing are governed by smart contracts. The vast majority of DeFi lending is over-collateralized loans secured by other crypto coins, which is primarily used for creating leveraged trading positions.

A typical transaction involves borrowing some of the stablecoins and putting Ethereum or Bitcoin as a collateral. Since the value of Ethereum or Bitcoin fluctuates there is a danger that the value of collateral can be lower than the borrowed amount. To mitigate this risk a smart contract uses an Oracle to obtain up-to-date cryptocurrency prices and automatically liquidate the position if the loan-to-value falls below a specified threshold. The threshold depends on the perceived riskiness of the collateral token and ranges between 50% and 80%.

A borrower has to pay a borrowing interest rate and can receive a lending rate on their collateral. In addition, a protocol collects a fee for its service, which goes to the pool controlled by protocol token holders. The lending rate is a function of the borrowing rate and the utilization of funds: borrowing fees, net of protocol fees, are spread among all lenders. The borrowing rate depends on the asset. It is set by the smart contract to maximize utilization of funds and changes in response to the market conditions.

Figure 5 shows that similar to DEX, the lending space is dominated by a few large

<sup>&</sup>lt;sup>16</sup>See Aoyagi (2020), Aoyagi and Ito (2021), Lehar and Parlour (2021), and Capponi and Jia (2021) for further results and comparison of decentralized and centralized exchanges.

players such as Aave, Anchor, and Compound protocols. Most protocols operate on a few chains. For example, Aave is built on three smart-contract platforms: Ethereum, Avalanche, and Polygon; Anchor only uses Terra, and Compound only Ethereum. Thus, the concentration within a particular smart contract platform is even higher.

#### [Fig. 5 About Here]

Figure A.6 in the Appendix shows the aggregated amount deposited and borrowed across different crypto tokens. The main activity is concentrated in stablecoins along with Ethereum and Wrapped Bitcoin (WBTC).<sup>17</sup> A large imbalance between the amount deposited and borrowed for Ethereum and Bitcoin means that investors use them as a collateral to borrow stablecoins, which can be used, for example, to buy Ethereum and Bitcoin, thus creating a leveraged position.

**Yield farming.** The desire to earn supersized returns led to the proliferation of smart contracts that aim to maximize the yield from holding crypto tokens. As we showed above crypto investors have several strategies to earn return on their coins. First, they can delegate their coins to validators who stake the coins and earn rewards for verification of transactions. Second, investors can earn fees for providing liquidity to DEXs. Third, they can earn an interest by depositing their coins into lending protocols. Finally, some token providers use airdrops — the practice of giving away tokens to a subset of investors meeting particular criteria.

The return on any of the above strategies varies over time. Yield farming smart contracts (or simply yield farms) aim to optimize the return by optimally allocating investments among multiple protocols and DeFi applications. The process also usually involves high leverage. For example, LP tokens obtained after placing tokens in a liquidity pool can be further used as a collateral or deposited into lending protocols.

The high leverage creates a risk of large losses due to a chain reaction of multiple contracts being liquidated when some contracts lose their value either during downturn market movements or because of some hacks. Also, while yield farm strategies are designed to maximize the yield on investment it does not automatically mean high returns because the underlying crypto tokens can loose value. In many cases, high yields are financed through in an increase in the token supply where the net effect depends on the investors willingness to absorb an ever increasing supply of tokens.

<sup>&</sup>lt;sup>17</sup>Wrapped Bitcoin is an Ethereum token that is intended to represent Bitcoin on the Ethereum blockchain. It is backed on a 1:1 basis with Bitcoin.

#### 3.2.3. NFT

Lastly, 2021 have seen a meteoric rise in hype and value of non-fungible tokens (NFTs). NFT is a unique piece of data stored on a blockchain. The data can be associated with a particular digital or physical asset or a license to use the asset for a specified purpose. Because each token is uniquely identifiable, NFTs differ from other cryptocurrencies. NFTs can be bought and sold, and are seen as a form of digital art. The space attracted attention in March 2021 when a digital collage of 5,000 images by the artist known as Beeple was sold at an eye-popping price of \$69 million at the Christie's auction house. The combined value of all NFTs in the end of January 2022 stood at about \$13 billion.<sup>18</sup>

## 4. DeFi vs. traditional financial system

Many of the existing problems with intermediaries originate from well-known economic frictions that are inherent in financial markets, such as asymmetric information, adverse selection, moral hazard, etc. This creates opportunities for abuse and also significant costs of guarding the public and the economy against financial fraud, malfeasance and systemic risk. Technological innovations have a long history in finance in helping to provide solutions to the above problems and improving the efficiency of financial markets.

DeFi applications thus far have had limited scope; they have been mainly built around simple applications such as trading in cryptocurrencies or collateralized lending. But they are growing rapidly in scope and complexity. They also escaped the burden of regulation and consumer protections and benefited from tremendous investor optimism that allowed many problems and inefficiencies to go unnoticed.

In what follows, we aim to highlight the important trade-offs offered by the two architectures. When comparing the potential benefits of DeFi solutions with those offered by the traditional system, it is important to think about the proposed new solutions in the context of the larger financial architecture rather than narrowly focusing on individual dimensions of possible inefficiencies.

## 4.1. Data privacy and transparency

How to protect data privacy in an increasingly digital society has become a major concern to regulators, activists and regular citizens alike. Crypto enthusiasts often tout the anonymity of transactions "a feature, not a bug" and view it as a major

<sup>&</sup>lt;sup>18</sup>https://nftgo.io/overview.

benefit over the traditional model, where the failure or corruption of a centralized intermediary could lead client data to be mistakenly exposed or hacked. While it is in the commercial interest of intermediaries to protect the privacy of their clients, it is a reasonable concern that intermediaries might not endogenize the full cost to the clients.<sup>19</sup> This conflict leads to a classic underinvestment problem relative to what consumers would prefer. In addition, financial intermediaries might have an interest in using client data for their own commercial purposes, or allowing third parties access, including the government.

Recognizing this problem, in the US a large set of regulations, such as the Bank Secrecy Act, Right to Financial Privacy Act, the Gramm-Leach-Bliley Act, and the Fair Credit Reporting Act, are put in place to protect consumers from unlawful access to their financial accounts by private and public institutions, and the unlawful disclosure or commercial use of financial information.

But the laws also recognize an important trade-off between individual privacy and other important societal goals such as preventing malevolent actors from using the financial system for money laundering, financing of criminal and terrorist activities, or tax evasion. This is typically achieved by putting into place Know-Your-Customer (KYC) and anti-money laundering (AML) laws that require financial institutions to verify the identity of a client when opening an account and to provide government authorities with information about suspicious financial transactions. Financial intermediaries in the traditional system then play the dual role of acting on the one hand as a "shield" to prevent the unauthorized collection, use, and disclosure of sensitive data. But on the other hand, they selectively grant access to information in well-defined circumstances where access to such data is important for the functioning of the economy or the broader society. Examples include reporting of capital gain tax to the IRS, or granting access to financial accounts of individuals in cases where an illegal or terrorist intent has been clearly defined by law and regulation.

Cryptocurrencies built on permissionless protocols preserve privacy by design by not collecting any personal information about account holders. Crypto tokens are represented by alphanumeric strings and protected by cryptography algorithms. Crypto addresses are very easy to generate and many protocols encourage users not to use an address more than once. Even if a protocol has a complete record of transactions, the identity of the person behind the transactions cannot be established unless this person uses the tokens to transact with an entity that does enforce KYC norms, such as regulated financial institutions. In many ways, the current modus operandi

<sup>&</sup>lt;sup>19</sup>Some infamous recent examples of data breaches in the financial are the 2017 breach of Equifax that exposed personal information of 147 million people or banks like Capital One and First American Financial Corporation, see https://www.upguard.com/blog/blogest-data-breaches.

of cryptocurrencies is similar to an old Swiss model of banking where people could set up anonymous accounts and no questions were asked. This model, however, has been rejected in the majority of developed countries in favor of more transparency and accountability.

Collecting and protecting data is not costless, and in the traditional architecture, intermediaries bear this cost. The benefits of relying on intermediaries as the important "entry nodes" for participants in the traditional financial system, means that KYC norms or anti-money laundering laws have to be monitored only at a limited set of nodes. For example, when a customer makes a payment using a credit card or a bank transfer from a US bank, a retailer does not need to worry about the legality of the funds. Similarly, the ability to collect taxes depends on the government's capacity to trace transactions and link them back to a person or organization. In the traditional system, centralized intermediaries such as exchanges or brokers are responsible for reporting transactions to the IRS.

The permissionless and pseudonymous architecture of DeFi generates formidable challenges for tax enforcement, aggravates issues of money laundering and other kinds of financial malfeasance, and as a result creates externalities on the rest of the economy. If entry into the system is not monitored by intermediaries but happens completely anonymously by setting up an address on a blockchain, KYC norms and AML laws would need to be regulated at the transactional level. In many cases this could be prohibitively costly, or impractical, and therefore lead to an in-transparent environment that facilitates illegal transactions.

Consider, for example, trading on a decentralized exchange. Recall that a decentralized exchange is simply a smart contract that executes trading between any pair of cryptocurrencies and that can be deployed anonymously by anyone. Suppose a customer trades and realizes some capital gains. Since the identity of the person behind the transactions cannot be established until this person uses the tokens at an entity that does an ID check, by transacting with entities that do not verify any ID, the person could spend the tokens linked to the capital gains transactions, and thus avoid ever paying capital gains taxes.

But even if the person transacts with an entity that does enforce KYC standards, it does not reveal the capital gains tax that is associated with the past transactions of this coin. In order to impute the true capital gains tax, the entity would need to either investigate the full history of transactions up to the current point or would need to delegate this task to another intermediary. In practice, tracing transactions along often multiple protocols is a challenging problem. Specialized blockchain analytics companies such as Bitfury Crystal and Chainalysis have shown that it can be successfully done in select cases of illegal transactions. However, successfully tracing all transactions will likely be very costly. Makarov and Schoar (2021) show for example, that Bitcoin flowing out of dark net markets like Hydra can be laundered through many intermediary addresses and can eventually enter KYC compliant exchanges such as Coinbase or Gemini without being tagged.

The pseudonymous nature of cryptocurrencies also makes it much harder to enforce rules against market manipulation, insider trading, and self-dealing, since suspicious transactions cannot easily be traced back to individuals. For example, large holders of cryptocurrencies have strong incentives to lobby government officials or regulators to promote investments in cryptocurrencies and adopt lax regulation. Especially at the early stages in the development of new technologies, any announcements endorsing the official use of cryptocurrencies create significant positive price impact, see Auer and Claessens (2020). The danger is that some regulators or politicians (or their friends) receive gifts in the form of cryptocurrencies (or simply already own cryptocurrencies) which would tilt their decision towards adoption even if it is not in the interest of the general public.

As the above discussion shows, to safeguard society against these inherent risks, a completely new framework of ensuring KYC and AML standards would have to be developed. The majority of DeFi players actively lobby that they should not be bearing the costs of linking transactions to economic actors and ensuring that the financial system preserves an adequate level of transparency and accountability citing technological constraints or the danger of losing a competitive advantage in the crypto space.<sup>20</sup> But unless society gives up entirely on collecting taxes and implementing KYC and AML practices, somebody has to bear these costs.

## 4.2. Economic rents

Another important dimension by which to assess a financial system is how economic rents are distributed among agents in the system. An important concern with the traditional financial system has been that the centralized position of intermediaries can allow them to extract excess economic rents at the expense of their customers. The proponents of the DeFi architecture typically argue that the open-source and permissionless nature of DeFi protocols promotes competition. Therefore, the claim is that DeFi solutions should drive out excess rents.

This view, however, neglects the fact that free entry is not synonymous with more

<sup>&</sup>lt;sup>20</sup>See, for example, https://stakingfac.medium.com/staking-infrastructure-providersunite-in-the-european-blockchain-association-6ecebbb8139f, https://republicansfinancialservices.house.gov/news/documentsingle.aspx?DocumentID=408238.

competition and thus not a panacea for beneficial outcomes in many situations. The effectiveness of competition depends on a number of factors such as whether there are barriers to entry, switching costs, product differentiation, asymmetric information, and network externalities. The presence of any of these factors hinders competition, and in some cases, even creates adverse effects from competition. Technological changes that affect any of these factors, therefore, also transform the competitive landscape.

Similar to the traditional financial system, there are several natural points where rents can accumulate at different layers in the DeFi architecture due to endogenous constraints to competition.

First, at the level of validators of transactions, in both PoW and PoS rents can accumulate due to inherent economies of scale and scope. In theory, in PoW protocols if miners were fully decentralized one could expect them to earn zero rent in steady-state because of free entry. In practice, however, as we showed in Section 1.1, mining is concentrated in pools and at the level of individual miners. High concentration of mining power can facilitate collusion and help sustain transaction fees above their average costs. For a dominant protocol such as Bitcoin, the competition from other PoW protocols can be limited because of the negative externalities the dominant network has on the security of smaller PoW networks. In particular, mining capacity can be redirected to launch 51% attacks on the smaller networks, as discussed in section Section 1.1.

Similarly, rents can also accrue to validators in PoS protocols. We showed in Section 1.2 that validators in PoS are concentrated. Furthermore, the same validators are active across a large cross-section of cryptocurrencies, effectively forming a new market structure. These validators control a large proportion of wealth that gives them substantial competitive advantage over newcomers with small amount of wealth.

Second, rents can also accrue at the level of the smart contract platforms that are built on the base layers. Similar to traditional payment systems like Visa, Mastercard or Paypal there are strong network externalities. Smart contract platforms differentiate themselves by the choice of programming language to code up smart contracts and the network architecture, and often have limited degree of interoperability. While smart contracts built on the same protocol can interact seamlessly with each other, communication between applications built on different platforms in general is limited.<sup>21</sup>

Naturally, the decision on which platform to build an application on depends on the existing pool of applications already deployed on the platform and the platform's future growth prospects. A popular platform with a wide range of applications and

 $<sup>^{21}</sup>$ A number of solutions have been proposed and are being developed to increase interoperability between chains, see for example, https://ethereum.org/en/bridges/ for more details.

a large user base provide better business prospects, and therefore, is more attractive than a less popular platform. Often these network effects increase exponentially with each user. As a result, developers and users might choose a more popular platform even if it charges higher transaction fees. These network externalities might also stand in the way of switching to a platform with a better technology if a critical mass of users is captured by the incumbent platform.

One could argue that even if the platform is a monopolist, competition between validators on that platform will keep fees low. However, as we showed above high concentration of validators can lead to collusion and allow them to earn excess rents. Even if validators do not collude, high transaction fees can still be realized if the platform operating capacity is limited and users need to pay a premium for priority execution, Huberman et al. (2021). Finally, the majority of PoS protocols have a minimum level of transaction fees as a protocol parameter, which provides the platform with a direct tool to limit competition among validators and earn rent.

Figure A.7 in the Appendix shows total transaction fees in the year 2021 across different platforms. The case of Ethereum is striking. The platform generated \$10 billion in fees from about 460 million transaction. In contrast, Visa's total revenue was \$24 billion over 165 billion transactions.<sup>22</sup> Thus, an average Ethereum fee per transaction has been 100X that of Visa.

For PoS platforms, an added complexity arises from the fact that the going concern value of the platform also affects the security of the platform itself and the applications that run on it. Since the value of the platform depends on the level of transaction fees, fees should be high enough to deter possible attacks on the platform, which can further support the platform's rent in equilibrium. These security concerns can also decrease competition among platforms. Since a low-value platform can be more easily attacked, the concerns over the platform's security may lead to slower growth, which in turn, can reduce the platform's current value.

Third, economies of scale at the level of individual DeFi applications can allow them to assemble local monopoly power and extract rents despite the open-source architecture of the blockchain. In addition, while in theory crypto smart contracts are usually described as an open-source code, in practice successful applications have tried to protect its code and limit its distribution. Here, an example of two decentralized exchanges, Uniswap and SushiSwap, is instructive.

Originally, Uniswap V2 was operated as an open-source software utilizing a General Public License (GPL), which allows anyone to run, distribute, and/or modify its code.

 $<sup>^{22}</sup>$ This figure is larger than transaction fees alone since Visa earns revenue from sources other than fees paid by direct users.

This has been used by a pseudonymous developer called Chef Nomi to create a clone of Uniswap called SushiSwap. Similar to centralized exchanges, decentralized exchanges are subject to economies of scale. An exchange with a large liquidity pool is preferred over an exchange with a small one. Therefore, an exchange clone will typically find it difficult to challenge the original exchange.

To compete with Uniswap, SushiSwap introduced a new business model, which is now adopted by a majority of other applications. The main change made by Chef Nomi was to create a governance token (SUSHI) and give it as a reward to traders who provide liquidity to the platform. The token allows its holders to vote on how the SushiSwap platform is run and potentially receive a portion of the transaction fees. As a consequence, investors can trade these tokens and speculate on the future prospects of the platform. This business model strengthens network externalities, and therefore, limits copycat strategies and competition. The more valuable the platform and its tokens are, the higher is the reward for liquidity providers. A larger liquidity pool, in turn, attracts more trading on the platform, which makes the platform more valuable.

The SUSHI token was also used to launch a "vampire attack" to drain liquidity out of Uniswap, whereby SUSHI tokens could be exchanged for Uniswap liquidity pool (LP) tokens. Those LP tokens would then be exchanged for the original assets put into the Uniswap liquidity pools, thus creating liquidity for SushiSwap instead. The attack was successful, draining Uniswap of about 55% of its liquidity.<sup>23</sup>

In response, Uniswap introduced its own governance token (UNI). To limit copycat attacks, the new version of the protocol, Uniswap V3, also adopted a different license agreement called Business Source License (BSL), which incorporates copyright law and allows Uniswap governance to restrict unauthorized commercialization of an entity's source code for two years.

Finally, rent extraction can even be driven by frictions at the customer level due to lack of financial literacy or behavioral biases. Many financial products today including smart contracts are complex contracts with multiple features. If consumers lack the financial sophistication to understand these product features, institutions that issue these contracts can shroud the actual cost of a product or service. A typical shrouding technique is to advertise or draw attention to one set of attractive features but hide other more expensive ones. If consumer are unable to analyze what is the best product even competition might not prevent rent extraction. In fact, more competition might lead to more shrouding as competing firms try to appeal to consumers with evermore enticing and salient features while hiding the unappealing dimensions of the product.

 $<sup>^{23} \</sup>tt https://www.vklaw.com/ImagineThatIPLawBlog/uniswap-v3-employs-a-new-license-agreement.$ 

Consumer finance products are often designed and marketed in this fashion, which leads to differential targeting of customers based on their financial literacy, see for example Célérier and Vallée (2017) or Ru and Schoar (2016). Similarly, in the crypto space, practices such as airdrops, yield farming, and meme DeFi tokens have helped capture interest of many investors, but many industry insiders question their value.<sup>24</sup>

### 4.3. Transaction costs

Even if a financial system limits economic rents it can still be inefficient because of high transaction costs. The traditional financial system has many inefficiencies, which result in high cost of banking services and long settlement time of transactions. A substantial part of these costs comes from need to cover brick-and-mortar costs of traditional banks and outdated infrastructure. Many banks today still use customized software that goes back to 1980s that lack real-time account reconciliation and liquidity management capabilities.

While many technological advances are largely exogenous to banks' actions the decision to when and how to implement them depends on the financial architecture. Centralized intermediaries can have limited incentives to invest in new technologies that could threaten their centralized position even if they are welfare improving. Also, modernizing bank's internal system can have a limited effect if other banks do not co-ordinate on the change. Often the threat of losing business to new entrants is necessary to force the incumbents to adopt more efficient technology.

The development of blockchain technology has had certainly positive effect on the incentives of financial industry to upgrade its infrastructure and reduce costs. It is less clear, however, to what extent the potential to reduce the costs depends on the permissionless nature of blockchain. In many cases, arguments can be made that a permissioned blockchain could designed to deliver a more cost-efficient and robust solution without curtailing competition.<sup>25</sup>

Notice also that the permissionless and open-source nature of a protocol does not necessary make an innovation process easy. It is often argued that if a blockchain protocol is inefficient then one can create an improved version (aka hard fork) by copying and upgrading the existing code. We showed in Section 4.2 that competition can be limited between different protocols because of strong network externalities and miners/validators can earn rent in equilibrium. If a new fork leaves less rent to min-

<sup>&</sup>lt;sup>24</sup>https://decrypt.co/40939/defi-meme-coins-no-joke-billions-ethereum,

https://decrypt.co/40318/vitalik-buterin-ethereum-defi-yield-farmers-central-banks. <sup>25</sup>See e.g., https://www.swift.com/news-events/news/swift-completes-landmark-dlt-proof-concept.

ers/validators they can have limited incentives to support it. Bier (2021) details the fight among Bitcoin developers about the Bitcoin protocol parameters that occurred in 2015-2017 and provides additional insights into challenges that come with forking a competing blockchain.

## 4.4. Governance

The promoters of cryptocurrencies often highlight the idea that the blockchain ledger removes the need for a trusted third party in the execution of contracts. However, this does not mean that the system can function completely devoid of any human intervention. Even if the execution of transactions and smart contracts on the blockchain are automated, the rules governing the blockchain itself and any upgrades to the system must be agreed upon and implemented by its participants. These rules define the governance of the system and in turn how it represents the interest of its different stake holders.

The major stakeholders in a blockchain eco-system are first the core developers who are charged with writing and updating the code that runs the blockchain. The validators who verify transactions and ensure the integrity of the blockchain are the second set of stake holders. Often, they decide if they want to adopt the changes provided by the developers. The third important group are the token holders. We can think of these as investors or equity holders. Finally, the forth group are users of the platform. On some platforms, the third and the fourth groups are the same people.

While all stakeholders have an interest of making the cryptocurrency they are engaged with succeed and grow, their incentives are not always completely aligned. For example, the users and developers might want fees on the blockchain to be low to make utilization more attractive, while investors and validators want to maximize the return on their financial investments. Stakeholders might also differ in their non-pecuniary benefits, e.g. some participants might be willing to forgo economic benefits for other objectives such as maintaining the independence or purity of the blockchain or possibly to undermine other blockchains as discussed before.

Thus, the classic problems in governance apply also to the crypto universe: Rules have to be set to facilitate coordination and provide incentives to adopt value increasing investments; and to prevent minority stakeholders from being expropriated by powerful insiders. Providers of capital are particularly prone to expropriation, since once the investment is made they do not have continued value added or recourse to the firm.

Corporate governance has been a prominent issue probably as long as organizations exist; in academic research the topic has attracted an enormous body of research at least since the publication of Berle and Means' famous book in 1932.<sup>26</sup> While there is significant heterogeneity across countries in the specific corporate governance rules, academic research has shown that private solutions even in competitive financial markets cannot generally resolve governance issues, and the recourse to the legal system is a crucial prerequisite for a well functioning financial system, see for example La Porta et al. (2000).

But this reliance on legal enforcement clashes with the maxim of decentralized finance that tries to avoid placing trust in any actor or institution, including the legal ones. In response to this challenge, decentralized finance has tried to develop a new form of governance, so-called decentralized autonomous organization (DAO). The basic idea of DAO is to spread control over decisions among all interested stakeholders. This is done by issuing special "governance" tokens that give their holders the power to propose changes to the protocol and vote on them. All activity is governed by smart contracts and recorded on the blockchain. In most DeFi applications one governance token equals a vote and new proposals are implemented according to a predefined majority rule. To ensure that the holders of governance tokens have an interest in the success of the platform long term, protocols often channel a share of the network's transaction fees into the wallets of the governance token holders. The tokens may also carry non-governance rights, like the right to be exchanged for certain other tokens at predefined rates. A famous example of DAO is MakerDAO.<sup>27</sup> Here is how DAO is explained on the Ethereum website:<sup>28</sup>

"Starting an organization with someone that involves funding and money requires a lot of trust in the people you are working with. But it is hard to trust someone you have only ever interacted with on the Internet. With DAOs you do not need to trust anyone else in the group, just the DAO's code, which is 100% transparent and verifiable by anyone."

But while a transparent and verifiable governance process is certainly an important first step, it does not necessarily ensure good governance. Any DAO design faces the same fundamental trade-offs and issues as traditional organizations. First, decision making in a fully decentralized organization can be inefficient. When the ownership is dispersed and stakes are small no owner might find it in their interest to spend effort and invest in learning about all the complexities needed to make a decision. As a result, many stakeholders might refrain from voting or lend their votes to a party that is trying to amass voting rights for self-interested reasons. Second, there is always a danger that investors with large stakes (blockholders) can capture the control and impose their

<sup>&</sup>lt;sup>26</sup>For an overview see Hermalin and Weisbach (2017).

<sup>&</sup>lt;sup>27</sup>https://makerdao.com/en/governance.

<sup>&</sup>lt;sup>28</sup>https://ethereum.org/en/dao/.

preference on the system. Recognizing this problem, corporate laws usually impose strict disclosure rules on blockholders. Emulating similar rules on a public permissionless blockchain would be challenging since everyone can control multiple anonymous accounts. Third, the voting system can give more power to participants, who may only be interested in maximizing short term profits as opposed to developing the protocol towards innovative use cases.<sup>29</sup> These arguments are very similar to the debate about investor short-termism in traditional governance, see Roe (2020).

Not surprisingly, the crypto space is abundant with colorful examples of governance issues.<sup>30</sup> Ultimately, the majority of insiders recognizes the inherent tensions posed by greater decentralization. Figure 6 shows that in the majority of crypto projects, developers and early investors chose to keep control of the platform by allocating significant stakes to themselves. In addition, even if developers do not have a large stake, in many cases they managed to maintain de facto significant control over the platform, e.g., Vitalik Buterin who has been dubbed the "benevolent dictator for life."<sup>31</sup>

#### [Fig. 6 About Here]

There has been little evidence so far to suggest that the crypto space can successfully resolve governance issues without relying on some off-chain mechanisms. Given that governance issues of blockchain platforms and traditional financial firms are not materially different, it is very likely that robust governance mechanisms will require the support of external regulation.

The history of corporate governance demonstrates that simply providing incentives for managers or investors, might not be sufficient to deter bad actors if the financial gains from misbehavior are large. As the implementation of governance rules in the US has shown, the ability for personal accountability of managers and directors is centrally important, see Bartlett and Talley (2017). Fiduciary duties that hold corporate agents personally accountable play a critical role in the enforcement of governance rules. The idea is that the threat of punishment creates disincentives from fraudulent behavior, where just losing some money from fraud would not have the same effect.

The pseudonymous nature of the permissionless blockchain environment, however, can make it difficult to hold bad actors accountable for their actions in the same way.

<sup>&</sup>lt;sup>29</sup>In fact, many recent attacks on DeFi apps exploited the possibility of taking over the voting mechanism to divert funds to the attacker, see for example https://www.wsj.com/articles/crypto-thieves-get-bolder-by-the-heist-stealing-record-amounts-11650582598.

<sup>&</sup>lt;sup>30</sup>See, for example, Bier (2021) and an attempted hostile takeover of Steem: https://decrypt.co/ 38050/steem-steemit-tron-justin-sun-cryptocurrency-war.

<sup>&</sup>lt;sup>31</sup>https://en.wikipedia.org/wiki/Benevolent\_dictator\_for\_life. Vitalik Buterin has also been one of the prominent critics of the DAO. See, for example, https://vitalik.ca/general/2021/08/16/voting3.html.

## 4.5. Systemic risk

One of the main sources of systemic risk in the traditional financial system is the reliance on fractional reserve banking. When banks take deposits from the public they only need to hold a fraction of these deposits in liquid assets as a reserve and can lend the remainder out to borrowers. The goal of a fractional reserve system is to expand economic activities in the economy by freeing capital for lending. It permits banks to use the majority of the deposits to generate returns in the form of interest rates on loans. The efficiency, however, comes at a cost of possible bank failures and runs.

DeFi so far has been operating under a narrow banking model, where every loan is overcollateralized.<sup>32</sup> Narrow banking removes many of the problems faced by fractional reserve systems, but it also constrains the efficient use of capital. The main risk comes from the ability of investors to take leveraged positions and a potential run on stablecoins.

A run on stablecoins can occur for a number of reasons. For stablecoins backed by traditional assets a run can happen for similar reasons to a run on bank or money market funds. In the absence of timely information about reserves, if investors doubt the quality of the collateral, they have an incentive to exchange the stablecoin for cash, causing a run unless the stablecoin is backed one-for-one with liquid assets like cash or short-term Treasuries. Possible solutions range from issuing stablecoins through insured banks, requiring stablecoins to be backed one-for-one with safe assets, to establishing a central bank digital currency. For a detailed discussion of the regulation of stablecoins and the trade-offs between private and central bank digital currencies, see Gorton and Zhang (2021) and Gorton (2021).

The situation is more complex in the case of algorithmic stablecoins that rely on intricate designs meant to help maintain the peg. Here the issue is less about transparency or misreporting because the design of a stablecoin is public knowledge and all transactions are recorded on the blockchain. Rather the main concern is about the complexity and potential fragility of the system. Since algorithmic stablecoins are not fully backed by safe assets, it is reasonable to expect that irrespective of a particular design, there always exist states of the world where the peg is broken and there can be a run on the stablecoin. The available documentation of stablecoins universally lacks rigorous analysis and contains only claims that the design is robust, which as the case of Iron Finance's Titanium proves can just be wishful thinking.<sup>33</sup>

<sup>&</sup>lt;sup>32</sup>There have been isolated examples of undercollateralized loans. See https://medium.com/ coinmonks/the-current-state-of-undercollateralized-defi-lending-2021-1f84e14527b5 for an overview of the current solutions.

<sup>&</sup>lt;sup>33</sup>https://forkast.news/iron-finances-defi-bank-run-why-mark-cuban-got-rekted/.

The ability to establish highly leveraged positions is another source of systemic risk. The crypto eco-system is famous for its wide range of highly leveraged products, with many exchanges offering up to 100X leverage for perpetual derivative contracts.<sup>34</sup> Figure A.8 in the Appendix shows that starting from July 2021 volume in crypto derivatives dominated the volume in spot markets. High leverage exacerbates volatility and as many industry observers believe is responsible for strong de-leveraging cycles and associated sharp drops in the cryptocurrency prices.<sup>35</sup>

DeFi adds an additional complication to the picture. Many DeFi protocols facilitate leverage and accept other protocols' tokens as collateral. Even though every smart contract and transaction are recorded on a blockchain, and therefore, in theory, could be analyzed, in practice multiple interconnecting contracts interacting with pseudonymous accounts result in a highly complex, and potentially fragile system. This fragility could potentially be acerbated if some critical smart contracts have unintended coding bugs.

## 5. Regulation

As discussed several times before in this article, the new financial architecture proposed by cryptocurrencies and decentralized finance, presents formidable challenges for regulators. Regulations of financial assets and services typically have three broad goals: (1) prevent the use of funds for illicit activities, money laundering or tax evasion; (2) protect participants in financial markets against fraud and abuses; (3) and ensure the integrity of markets and payment systems and overall financial stability.

Our discussion in Section 4 highlights that at present DeFi solutions do not comply with the above three goals. If society does not want to give up on ensuring these goals, some form of technological and regulatory solution seems desirable. In the response to the rise of cryptocurrencies, different countries have followed vastly different approaches. For example, China officially banned trading in cryptocurrencies and developed its own central bank digital currency, while El Salvador allowed Bitcoin as legal tender. In the US the regulatory environment is still in flux and there are overlapping responsibilities and sometimes even contradictory approaches.

However, there is urgency to provide a clearer regulatory framework for at least two reasons. First, regulatory certainty is always important for entrepreneurs and investors who wish to decide whether and how to participate in new technologies. Second,

<sup>&</sup>lt;sup>34</sup>See, e.g., https://www.bloomberg.com/news/articles/2022-01-26/wild-crypto-leverageon-offer-for-pros-in-20-times-bitcoin-bet.

<sup>&</sup>lt;sup>35</sup>https://insidebitcoins.com/news/how-crypto-derivatives-crashed-the-market-doprices-signal-buy-now, https://bitcoinmagazine.com/markets/leverage-in-derivatesmarket-dipped-bitcoin-price.

the exponential growth of cryptocurrencies can lead to a situation where the political economy of regulation becomes very difficult, if regulators wait too long. In effect, cryptocurrencies and DeFi applications can become too-big-to-regulate. We showed in Section 4 that currently many DeFi solutions do not bear the full cost of the externalities they impose on the economy, such as enforcing KYC or AML laws or complying with tax reporting. Part of the current valuation of some cryptocurrencies and DeFi applications might even be based on an expectation that they will not have to ever comply with these regulations. Thus, requiring DeFi solutions to start internalizing these costs will likely result in losses for some of the current investors. As a result, any such proposals are usually met with strong resistance by the crypto community. This puts regulators in a difficult position. While they need to safeguard the financial system, in a democracy there is often populist pressure to forgo doing the things that are good in the longer run to satisfy short run goals. In fact, the losses might be blamed on the regulatory action itself, rather than the attempt by regulators to prevent even larger losses for society going forward.

The main challenges in regulating cryptocurrencies stem from the pseudonymous and jurisdiction-free nature of this new architecture, which is a consequence of the use of permissionless blockchain protocols and the smart contracts running on them. The traditional financial architecture where access runs through centralized intermediaries allows each country to determine its own regulatory framework and decide for example, who can open a bank account, what documentation must be supplied, or how information can be collected and stored. Also, as the recent geopolitical experience between Russia and the West shows, the traditional system makes it possible to restrict the financial system of one country from accessing the financial system of other countries.

The anonymous and permissionless nature of DeFi apps and the underlying blockchain protocols have the potential to remove the boundaries between the financial systems of different countries or even enable citizens to transact in an eco-system that is completely outside of government regulation or tax enforcement. While financial integration can have benefits through better risk sharing or improved liquidity it can also have large costs, if poorly regulated systems undercut better regulated ones in a race to the bottom. This becomes especially prevalent if different financial systems operate with vastly different standards.

So what are the available options for regulators? While a complete discussion of all dimensions of regulation is beyond the scope of this paper, we outline a few key options for rule-makers. A natural place for regulatory oversight in this new eco-system is at the level of developers and validators, which in turn control the network protocol. Once this level of regulatory compliance is established, many other functions can be built. In

particular, separate entities can be established that would be responsible for verifying the identities and certifying that crypto addresses belong to confirmed users. These entities should be subject to regular audits. The protocols can be adjusted so that validators can check if a particular address belongs to a certified entity and validators would be charged with only processing transactions that involve certified addresses.

In addition, one could imagine that customers can also be provided with private keys based on their characteristics, such as financial wealth or sophistication. Smart contracts can be ranked based on their safety, risk, etc. Rules can be established that would allow different smart contract categories to interact with customers who can provide the required key. Smart contracts can be designed to automate the ranking of other smart contracts and automate the generation of private keys. Cryptography algorithms can be developed to guard customers privacy. Transitioning to this model will likely require some time and development of new solutions. Therefore, it would be important to lay out an appropriate timeline and deadlines so that market participants can prepare for a smooth transition.

Since countries might differ in how they want to structure their regulatory environment for validators, each country can opt to run its own version of the blockchain. But if some countries agree broadly on regulatory standards, they can use the same blockchain. Countries that choose to run separate versions of the blockchain can interact with others using interoperability mechanisms such as bridges. The above solution can be more easily applied to new blockchains. But, if a majority of large countries agree on coordinated regulation then even the existing blockchains can be brought into a legal framework without the need to break them up into separate sidechains based on different regulatory requirements.

The above solution looks similar to a permissioned blockchain, but this system preserves most of the desired properties of the original design of cryptocurrencies, e.g., transactions can be observable on the blockchain, settlement is immediate, and the same set of smart contracts can be executed on it. In addition, if many countries agree on regulation validators can be elected so that no country has a monopoly over the networks. The ability to regulate validators can potentially change the enforcement of smart contracts, by allowing recourse to the contracting parties. But as we discussed in Section 2 it can have a positive effect on efficiency.

In contrast, if regulators give up on the ability to oversee validators, the effectiveness of regulation will be much more limited and will depend on the goodwill and voluntary cooperation of validators and developers of the blockchain. If validators accept transactions from every party, the most regulators can hope for is to separate the network into "regulated" and "unregulated" parts. This could be done by requiring, say, US citizens to interact only with certified DeFi apps, which comply with KYC and AML regulations and provide reports on trades, tax compliance or other activities. The relative size of the regulated and unregulated networks will depend on the relative investment opportunities in these two networks and the ease of moving funds between them. The problem of regulating compliance only at the level of DApps is that first, many citizens even from countries that try to regulate DeFi applications, could still find it attractive to invest funds in the unregulated network to avoid paying taxes, etc. The ability to evade compliance, can provide a large subsidy for the unregulated part of DeFi apps. Second, since regulation will have generally a limited bite on the unregulated part it can harbor many bad actors and facilitate illegal activities. The opportunities of sidestepping the regulated part will generally increase with the level of crypto-adoption, since people will be able to predominately interact in the unregulated part and avoid triggering regulatory compliance.

## 6. Conclusion

In this article we provide an introduction of how the new DeFi architecture works and the mechanics behind it. We also lay out some of the potential benefits and challenges of the developing new system and present a comparison to the traditional system of financial intermediation. In our discussion we focus on the economic forces and frictions that can arise within this system and the regulatory approaches that might help to mitigate the problems. Our analysis highlights that while the DeFi architecture might have the potential to reduce transaction costs, it is not an automatic solution to the problem of rents in the financial sector. But it may also create additional problems. We identify as a key challenge to regulators, the permissionless and anonymous nature of the current DeFi blockchains. These provide the opportunity for market participants to circumvent controls in the financial system, and create externalities for the rest of society, for example through facilitating tax evasion or skirting of AML laws.

We highlight that there are ways to regulate the DeFi system which would preserve a majority of features of the blockchain architecture but support accountability and regulatory compliance. These solutions would rely on a system where validators on the blockchain agree to check if a particular address belongs to a certified entity and validators would be charged with only processing transactions that involve certified addresses.

How this system evolves, in terms of technology and regulation has important consequences for liquidity and credit provision to the economy, and ultimately the standing of the US and other global economies. There are also strategic and competitive implications across countries. The US obtains significant economic and strategic benefits from the central role that the dollar and the US financial system hold internationally. Therefore, it is in the US interest to encourage innovation and modern financial technologies but at the same time to set standards that protect consumers and maintain the transparency, accountability and stability of the system. The cross-jurisdictional structure of permissionless blockchain ledgers entails a danger that participants will engage in regulatory arbitrage which would undermine the financial system and its stability. Coordination between the main financial markets will be important to prevent a hollowing-out of financial regulations.

## References

- Allen, F., Carletti, E., and Gu, X. (2019). The Roles of Banks in Financial Systems. In Berger, A. N., Molyneux, P., and Wilson, J. O. S., editors, *The Oxford Handbook of Banking*, Oxford Handbooks. Oxford University Press, 3rd edition.
- Aoyagi, J. (2020). Liquidity Provision by Automated Market Makers. Available at SSRN 3674178.
- Aoyagi, J. and Ito, Y. (2021). Liquidity Implication of Constant Product Market Makers. Available at SSRN 3808755.
- Aramonte, S., Huang, W., and Schrimpf, A. (2021). DeFi risks and the decentralisation illusion. BIS Quarterly Review.
- Auer, R. and Claessens, S. (2020). Cryptocurrency Market Reactions to Regulatory News. Globalization Institute Working Paper 381, Federal Reserve Bank of Dallas.
- Bartlett, R. and Talley, E. (2017). Law and Corporate Governance. In *The Handbook* of the Economics of Corporate Governance, volume 1, pages 177–234. Elsevier.
- Beniiche, A. (2020). A Study of Blockchain Oracles. arXiv:2004.07140.
- Berg, T., Puri, M., and Rocholl, J. (2020). Loan Officer Incentives, Internal Rating Models, and Default Rates. *Review of Finance*, 24(3):529–578.
- Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.
- Bier, J. (2021). The Blocksize War: The Battle Over Who Controls Bitcoin's Protocol Rules. Independently Published.
- Boot, A. W. A. and Thakor, A. V. (2015). Financial System Architecture. *The Review* of Financial Studies, 10(3):693–733.
- Budish, E. (2018). The Economic Limits of Bitcoin and the Blockchain. Working Paper 24717, National Bureau of Economic Research.
- Caldarelli, G. and Ellul, J. (2021). The Blockchain Oracle Problem in Decentralized Finance—A Multivocal Approach. *Applied Sciences*, 11(16).
- Campbell, J. Y. (2016). Restoring rational choice: The challenge of consumer financial regulation. *American Economic Review*, 106(5):1–30.

- Capponi, A. and Jia, R. (2021). The Adoption of Blockchain-based Decentralized Exchanges. arXiv:2103.08842.
- Célérier, C. and Vallée, B. (2017). Catering to Investors Through Security Design: Headline Rate and Complexity. *The Quarterly Journal of Economics*, 132(3):1469– 1508.
- Cong, L. W. and He, Z. (2019). Blockchain Disruption and Smart Contracts. The Review of Financial Studies, 32(5):1754–1797.
- Cong, L. W., He, Z., and Li, J. (2020). Decentralized Mining in Centralized Pools. The Review of Financial Studies, 34(3):1191–1235.
- Cong, L. W. and Xiao, Y. (2021). Categories and Functions of Crypto-Tokens. In Pompella, M. and Matousek, R., editors, *The Palgrave Handbook of FinTech and Blockchain*, pages 267–284. Springer.
- Ferreira, D., Li, J., and Nikolowa, R. (2019). Corporate Capture of Blockchain Governance. Working Paper 593, European Corporate Governance Institute (ECGI).
- Gorton, G. B. (2021). The Orkney Slew and Central Bank Digital Currencies. *Available at SSRN 3937323*.
- Gorton, G. B. and Zhang, J. (2021). Taming Wildcat Stablecoins. Available at SSRN 3888752.
- Halaburda, H., He, Z., and Li, J. (2021). An economic model of consensus on distributed ledgers. Working Paper 29515, National Bureau of Economic Research.
- Harvey, C. R. (2016). Cryptofinance. Available at SSRN 2438299.
- Harvey, C. R., Ramachandran, A., and Santoro, J. (2021). *DeFi and the Future of Finance*. John Wiley & Sons, Newark.
- Hermalin, B. and Weisbach, M. (2017). The Handbook of the Economics of Corporate Governance. Elsevier.
- Huberman, G., Leshno, J. D., and Moallemi, C. (2021). Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. *The Review of Economic Studies*, 88(6):3011–3040.
- Kogias, E. K., Jovanovic, P., Gailly, N., Khoffi, I., Gasser, L., and Ford, B. (2016). Enhancing Bitcoin Security and Performance with Strong Consistency via Collective

Signing. In Proceedings of the 25th USENIX Conference on Security Symposium, SEC'16, pages 279–296. USENIX Association.

- La Porta, R., de Silanes, F. L., Shleifer, A., and Vishny, R. (2000). Investor Protection and Corporate Governance. *Journal of Financial Economics*, 58(1):3–27.
- Laibson, D., Repetto, A., and Tobacman, J. (2007). Estimating Discount Functions with Consumption Choices over the Lifecycle. Working Paper 13314, National Bureau of Economic Research.
- Lehar, A. and Parlour, C. A. (2021). Decentralized Exchanges. Available at SSRN 3905316. Working paper, University of Calgary and University of California, Berkeley.
- Lusardi, A. and Mitchell, O. S. (2007). Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education. *Business Economics*, 42(1):35–44.
- Makarov, I. and Schoar, A. (2021). Blockchain Analysis of the Bitcoin Market. Working Paper 29396, National Bureau of Economic Research.
- Merton, R. C. (1995). A Functional Perspective of Financial Intermediation. Financial Management, 24(2):23–41.
- Platt, M., Sedlmeir, J., Platt, D., Tasca, P., Xu, J., Vadgama, N., and Ibañez, J. I. (2021). Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work. arXiv:2109.03667.
- Prasad, E. S. (2021). The Future of Money: How the Digital Revolution is Transforming Currencies and Finance. Harvard University Press, Cambridge, Massachusetts.
- Roe, M. J. (2020). Stock Market Short-Termism's Impact. Law Working Paper 426/2018, European Corporate Governance Institute (ECGI).
- Ru, H. and Schoar, A. (2016). Do Credit Card Companies Screen for Behavioral Biases? Working Paper 22360, National Bureau of Economic Research.
- Schär, F. (2021). Decentralized Finance: On Blockchain- and Smart Contract-Based Financial Markets. *Review*, 103(2):153–174.
- Werbach, K. (2018). The Blockchain and the New Architecture of Trust. The MIT Press.

- Werbach, K. and Cornell, N. (2017). Contracts Ex Machina. *Duke Law Journal*, 67:313–382.
- Wright, A. and De Filippi, P. (2015). Decentralized Blockchain Technology and the Rise of Lex Cryptographia. *Available at SSRN 2580664*.

## **Figures**



Figure 1: Average transaction fee and price for Bitcoin and Ethereum. This figure shows the daily transaction fees and closing prices for Bitcoin and Ethereum from January 2017 to February 2022. Daily closing prices are plotted on the left-axis and daily average transaction fees are plotted on the right axis. The figures are plotted in log-scale. Data source: Messari.io. Authors' calculations.



**Figure 2: Share of market capitalization by token categories.** This figure shows the share of market capitalization by seven categories of cryptocurrency tokens and coins (here we collectively refer to them as tokens) as of February 2022. Smart Contract Platform include tokens for platforms that host smart contracts on their own blockchains. Stablecoin refers to tokens that are pegged to a specific asset such as fiat currency. dApp includes tokens used for different decentralized application protocols. NFT refers to non-fungible tokens. Other refers to the rest of the cryptocurrency tokens that cannot be classified to the categories listed above. Data source: CoinGecko. Authors' calculations.



**Figure 3:** Share of market capitalization by stablecoin categories. These figures show the share of market capitalization by stablecoin categories as of February 2022. Panel A shows the share of stablecoins backed by traditional assets comparing to those backed by crypto-assets. Stablecoins backed by crypto-assets include those algorithmically backed by a particular cryptocurrency or by multiple tokens such as tokens in a liquidity pool. Panel B shows the share of top stablecoins backed by traditional assets. Panel C shows the share of top stablecoins backed by crypto-assets. Data source: CoinGecko. Authors' calculations.



**Figure 4:** Market capitalization and total value locked of decentralized finance. These figures show the market capitalization and total value locked for different categories of decentralized finance from January 2019 to February 2022. Trading refers to tokens used in decentralized exchanges including those for spot trade and derivative exchanges. Lending and borrowing refers to DeFi platforms where lenders add funds into liquidity pools in return for a regular interest rate from borrowers. Yield Farming includes yield aggregators and protocols that incentive people to deposit or lend out their tokens in exchange of rewards. Panel A shows the market capitalization of the tokens in billion USD and Panel B shows the total value locked in billion USD. Data source: market capitalization data comes from CoinGecko and total value locked data comes from Defi Llama. Authors' calculations.



**Figure 5: Market capitalization and total value locked of decentralized lending.** These figures show the market capitalization and total value locked for the top 20 lending protocols based on market capitalization from May 2020 to February 2022. Panel A shows the market capitalization of the tokens in billion USD and Panel B shows the total value locked in billion USD. Data source: market capitalization data comes from CoinGecko and total value locked data comes from Defi Llama. Authors' calculations.



Figure 6: Initial Coin Offering Insider Share. This figure shows the insider shares from top 50 tokens' ICOs. Insider share includes tokens to founding teams and developers, early investors such as seed investors, venture capital firms, and private sale investors, and associated entities which include companies that are related to the protocols or protocol founders. Insiders shares do not include shares that go into the community such as airdrops, grants, rewards, and tokens to public sale investors, and shares for the development of protocols, such as those going into foundations and reserves. Data source: Messari.io. Authors' calculations.

### Table 1: Concentration of validator stakes

This table reports the concentration of validator stakes for the top ten proof-of-stake smart contract platforms by market capitalization as of February 2022. Validator stakes include stakes provided by validator themselves and stakes delegated to validators. Column 2 shows the amount staked as a percentage of circulating supply. Column 3 and 4 show the amount staked for the top 10 and 50 validators as a share of total amount staked on the platform. The data exclude Ethereum since it is in a transition period. Data source: Stakingrewards.com and author's calculations.

Cryptogurrongy	Amount staked as a $(\%)$	Validator	Concentration
Cryptocurrency	of circulating supply	top 10	top 50
Solana	70%	23%	56%
Cardano	73%	30%	47%
Avalanche	97%	17%	57%
Terra	77%	36%	76%
Polkadot	57%	30%	56%
Cosmos Hub	63%	45%	87%
NEAR Protocol	61%	50%	96%
Polygon	34%	72%	99%
Fantom	54%	88%	100%
Tezos	76%	63%	96%

#### Table 2: Top validators

This table reports the top 15 PoS validators and their aggregate stakes in the top ten proof-of-stake smart contract platforms by market capitalization as of February 2022. Validator stakes include stakes provided by validator themselves and stakes delegated to validators. Column 2 shows the aggregate dollar amount staked. Column 3 shows validators market share. Data source: Stakingrewards.com and author's calculations.

Validator	Staked USD	Share
Everstake	2.8 B	2.20%
Binance Staking	$2.6 \mathrm{B}$	2.10%
Chorus One	1.6 B	1.30%
Dokia Capital	1.6 B	1.30%
Certus One	$1.5 \mathrm{B}$	1.20%
<b>Bison</b> Trails	$1.5 \mathrm{B}$	1.20%
Allnodes	$1.5 \mathrm{B}$	1.20%
InfStones	$1.5 \mathrm{B}$	1.20%
Kraken	1.4 B	1.10%
Staked	1.2  B	1.00%
P2P Validator	1.2  B	1.00%
Orion.Money	1.1 B	0.90%
<b>B</b> -Harvest	1.0 B	0.80%
Staking Facilities	1.0 B	0.80%
Figment	1.0 B	0.80%

# Appendix



Figures

**Figure A.1: Market capitalization of top non-smart contract cryptocurrencies.** These figures show the market capitalization of top non-smart contract cryptocurrencies and the rest from January 2019 to February 2022. The top seven cryptocurrencies include Bitcoin, Dogecoin, Litecoin, Bitcoin Cash, Dash, and Decdred. Panel A shows the market capitalization of the tokens in billion USD and the Panel B shows their corresponding percentages as a share of market capitalization for all cryptocurrencies. Data source: CoinGecko. Authors' calculations.



Figure A.2: Market capitalization of top smart contract platforms. These figures show the market capitalization of top five smart contract platforms and the rest from January 2019 to February 2022. The coins for the top five smart contract platforms are Ethereum, Binance Coin, Solana, Cardano, and Polkadot. Panel A shows the market capitalization of the tokens in billion USD and the Panel B shows their corresponding percentages as a share of market capitalization for all smart contract platforms. Data source: CoinGecko. Authors' calculations.



Figure A.3: Total value locked on top smart contract platforms. This figure shows the multi-chain total value locked by top smart contract platforms from January 2021 to February 2022. Data source: Defi Llama. Authors' calculations.



Figure A.4: Daily spot trade volume for centralized exchanges compared to decentralized exchanges. This figure shows the daily spot trade volume for centralized and decentralized exchanges from January 2020 to February 2022. The figure is plotted in log-scale. Data source: The Block. Authors' calculations.



Figure A.5: Exchange concentration. These figures show the top decentralized exchanges (DEX) and centralized exchanges' (CEX) monthly trading volume concentration in 2021. Data source: CoinGecko Yearly Report 2021.



Figure A.6: Tokens supplied and borrowed. This figure shows the aggregated deposit and borrowing of the top 15 tokens for the top three lending protocols in million USD: Aave, MakerDAO, and Compound as of February 25, 2022. Data source: protocol statistics. Authors' calculations. https://app.aave.com/markets, https://compound.finance/markets, https://daistats.com



Figure A.7: Platform transaction fees. This figure shows the total fees and revenues in 2021 for Level-1 blokchains and two payment networks: Visa and Stripe. Source: The Year in Ethereum 2021: https://stark.mirror.xyz/q30nsK7mvfGtTQ72nfoxLyEV51fY0qUfJIoKBx7BG11.



Figure A.8: Monthly Spot vs. Derivatives Volume. This figure shows the monthly spot and derivatives trade volume in USD from February 2020 to January 2022. Data source: Cryptocompare Exchange Review, January 2022.

# Tables

Coin	Date	Succeeded?	Loss in USD	After the Attack
Bitcoin SV (BSV)	Aug 5, 2021	No	\	The BSV team claimed that the attack was thwarted and all fraud- ulent chains identified.
Bitcoin SV (BSV)	Aug 3, 2021	Yes	Unknown	The BSV team recommended that node operators invalidate the fraudulent chain. The Bitcoin Association collected evidence of the illegal activity and its representatives worked with law enforcement authorities in affected jurisdictions.
Bitcoin SV (BSV)	Jun 24, 2021 - Jul 9, 2021	Yes	Unknown	Several crypto exchanges suspended BSV transactions, deposits and withdrawals. After the July 6 block reorganization, the BSV team discovered the malicious nature of the activity, then took mitigat- ing and preventative measures. The exchange Bitmart later claimed that the attacker had deposited "fake" BSV, traded them for other coins, and moved those coins to other exchanges. On July 23, Bit- mart filed for injunctive relief in the Court seeking to prevent fur- ther transfers and asking for other exchanges to freeze coins they received from the attacker. Binance announced a shut-down of its BSV mining pool scheduled for July 31.

## Table A.1: Blockchain 51% Attacks

Coin	Date	Succeeded?	Loss in USD	After the Attack
Verge (XVG)	Feb 15, 2021	No	\	Bittrex paused the XVG wallet. The Verge team said the attack was thwarted and failed.
Firo (FIRO, formerly Zcoin)	Jan 18, 2021	Yes	\$4.5 million but more than 70% of the FIRO was recovered	Exchanges paused deposits and withdrawals. The Firo team issued an emergency switch to temporarily disable Lelantus to prevent the attacker from anonymizing funds. They also released a hotfix as a preventive measure on Jan 21, asking all wallets and masternodes be upgraded. The price of FIRO dropped -16.51% on January 21. The Firo team locked the attacker's proceeds and suspected that this attack was not financially motivated. The Firo community voted to support reimbursing exchanges with the locked funds. The funds were returned to Binance. The Firo team expedited activa- tion of ChainLocks, a secondary validation layer, and deloyed it on Jan 28.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Aeternity (AE)	Dec 5, 2020 - Jan 8, 2021	Yes	more than \$5,000,000 but all the stolen AE were recovered later	Huobi timely paused all the AE deposits and withdrawals and alerted the AE team on Dec 7, shortly after the attacker broad- casted the new chain. Binance delisted AE on Dec 30. The AE team found that this attack targeted specific exchanges (OKEx, Huobi, Gate.IO and Binance) in their investigations. The AE com- munity members helped to mitigate the 51% attack by renting hashing power to mine in the community fork. On Jan 3, the Ae- ternity Community Telegram group was attacked. The AE team claimed that they thwarted the attacker's attempt to roll back ex- change transactions on Jan 8, and recovered the 29 million stolen AE tokens. The AE team also announced their plan about imple- menting Hyperchains, which are PoS systems that rely on existing PoW blockchains to prevent 51% attacks.
Bitcoin Cash ABC (BCHA, now eCash)	Nov 28, 2020	Yes	Unknown	Given that the attack was not financially motivated but for protest- ing a new miner tax, the unknown attackers could not sustain this attack. BCHA's price was not adversely affected by the attack.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Grin (GRIN)	Nov 7, 2020	Unknown	Unknown	The motivation for this attack remains unclear. The develop- ment team put a warning on its website for the sudden increase of hashrate which coincided with the Nicehash rate doubling outside of known pools. It also suggested extra confirmations on transactions. The price of GRIN remained relatively unchanged after the news of 51% threat broke.
Ethereum Classic (ETC)	Aug 29, 2020	Yes	Unknown	The series of attacks had no significant impact on the price of ETC. On Aug 31, the ETC team announced that they would pursue en- forcement and regulation of hash rental. On Sep 1, NiceHash ac- knowledged its hash-power rental platform may have facilitated the attacks. The ETC later implemented an Modified Exponential Sub- jective Scoring (MESS) solution to reduce the likelihood of future 51% attacks.
Ethereum Classic (ETC)	Aug 6, 2020	Yes	Unknown	Bitfly and Binance reported the block reorganization and halted ETC transactions, withdrawals, and deposits. The exchange OKEx said it would consider delisting ETC due to the network's severe lack of security. Coinbase extended deposit and withdrawal confir- mation times for ETC to roughly two weeks. The ETC team an- nounced a security plan on Aug 19.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Ethereum Classic (ETC)	Jul 29, 2020 - Aug 1, 2020	Yes	\$5,600,000	The blockchain analytics firm Bitquery reported investigations that debunked the ETC team's initial statement of no attack.
Bitcoin Gold (BTG)	Jul 1, 2020	No	/	Besides warnings, the BTG team privately supplied mining pools and exchanges with an updated version of the BTG network which has a checkpoint that automatically rejected the attacker's chain.
Bitcoin Gold (BTG)	Jan 23, 2020 - Jan 24, 2020	Yes	\$72,000	BTG's market price went up about 18 percent over 24 hours after news of the attack broke. In a white paper, the BTG team pro- posed a new soft fork approach, Cross-Chain Block Notarization Protocol, to prevent future 51% attacks.
Vertcoin (VTC)	Dec 1, 2019	No	\	The motivation for this attack remains unclear. Bittrex, possibly the original target of the attack, disabled its wallet before the reor- ganized blocks were published, thus prevented the potential double- spend. The VTC developer blamed Nicehash for their hashpower rental services.
Expanse (EXP)	Jul 29, 2019	Yes	\$12	This attack received little news coverage. Only a former researcher at the MIT Digital Currency Initiative disclosed it on github.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Litecoin Cash (LCC)	Jul 4, 2019 - Jul 7, 2019	Yes	\$5,500	This attack received little news coverage. Only a former researcher at the MIT Digital Currency Initiative disclosed it on github.
Ethereum Classic (ETC)	Jan 5, 2019 - Jan 7, 2019	Yes	\$1,100,000	The ETC team initially claimed there was no attack but later con- firmed it. Coinbase published a report on the attack and paused all ETC transactions, withdrawals and deposits. ETC had a near 10% depreciation on Jan 7. The blockchain security firm SlowMist found the attacker returned stolen funds to the YoBit and Gate.io exchanges on Jan 10.
Vertcoin (VTC)	Oct 12, 2018 - Dec 2, 2018	Yes	more than \$100,000	Coinbase published a report that provids many details on the time- line and financial losses of this series of attacks. The VTC developer blamed cloud-mining services such as Nicehash.
AurumCoin (AU)	Nov 9, 2018	Yes	\$550,000	The Aurum Coin team put all the blame on the exchange Cryp- topia, and claimed that the AU team is not responsible to the loss because AurumCoin is an open-source distributed crypto currency. Cryptopia did not even acknowledge the loss.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Pigeoncoin (PGN)	Sep 27, 2018	Yes	\$15,000	The PGN developers patched the bug that was exploited in the at- tack. Because PGN is a copycat crypto currency, the bug was orig- inally from the Bitcoin source code which was already fixed on Sep 19. Trading resumed on Oct 2.
Ravencoin (RVN)	Sep 13, 2018 - Sep 14, 2018	Yes	Unknown	The Ravencoin team reported their findings and solutions on Sep 18. They chose to implement a default maximum reorg depth with specific node conditions as a solution to prevent future 51% attacks. They also released a hotfix for a bug that was inherited from Bit- coin source code which allows double-spend attacks using the chain on Sep 21.
FLO Blockchain (FLO)	Sep 8, 2018	Yes	\$27,500	Bittrex disabled the wallet after the double-spend and alerted the FLO team. The FLO team decided to repay the approximate 700,000 FLO stolen from Bittrex and asked the FLO community for donations. To mitigate 51% attacks and protect the network, the FLO team initially planned to implement Sunny King's advanced checkpointing system, but later chose to add the more applicable max reorg depth consensus rules to FLO instead of using the cen- tral checkpoint mechanism.

Coin	Date	Succeeded?	Loss in USD	After the Attack
ZenCash (ZEN, now Horizen)	Jun 2, 2018	Yes	more than \$600,000	The ZenCash team announced that they had taken mitigating ac- tions, contacted exchanges to increase confirmation times, and con- ducted forensic analysis soon after receiving warning from a pool operator. On June 3, the Zen team released an official statement about the attack on their website.
Litecoin Cash (LCC)	May 30, 2018	Yes	Unknown	The exchange YoBit tweeted that a 51% attack on LCC was iden- tified. The LCC team alerted exchanges to increase confirmation requirements, and announced that there would possibly be a hard fork. Some news reports implied that the loss was minor in this at- tack. Later in a white paper, the LCC team proposed a new hybrid PoW/PoS solution, "The Hive", that aims to protect the network against 51% attackers.
Verge (XVG)	May 22, 2018	Yes	more than \$1,700,000	After attackers exploited the same weakness as the previous April attack, Verge tried to downplay it as a DDos attack on some mining pools. The price of XVG dropped significantly after the attack.
Bitcoin Gold (BTG)	May 16, 2018 - May 19, 2018	Yes	\$18,000,000	The BTG team updated its mining algorithm in June 2018 in order to add an immediate measure of safety from 51% attacks. Although the BTG team warned exchanges about the attack, the exchange Bittrex asked BTG team to pay for their loss. BTG refused to pay and was delisted from Bittrex later in September 2018.

Coin	Date	Succeeded?	Loss in USD	After the Attack
MonaCoin (MONA)	May 13, 2018 - May 15, 2018	Yes	\$90,000	Many exchanges halted deposits of Monacoin after the news of at- tack. The Monacoin developer advised exchanges to increase con- firmations to 100. Some news report stated that the attacker had been attempting to exploit a weakness in the Monacoin's difficulty adjustment mechanism for six months prior to this attack being de- tected.
Verge (XVG)	Apr 4, 2018	Yes	\$15,000	The problem was temporarily fixed with an emergency commit posted by the lead Verge developer, because the attackers used a weakness in the Verge code to falsify time stamps on blocks. Critics said that the vulnerability remains unfixed after the blockchain was hard-forked. The Verge team tried to downplay the severity of the attack on social medias.
Elec- troneum (ETN)	Apr 1, 2018	Yes	Unknown	It was first noticed because a massive amount of empty blocks were constantly mined on the currency's blockchain. Some ETN commu- nity members suspected that the attacker was Bitmain, who seemed to have large proportion of network hashrate at that time. This at- tack affected the Electroneum for a while, but Electroneum eventu- ally moved on.

Coin	Date	Succeeded?	Loss in USD	After the Attack
Krypton (KR)	Aug 26, 2016	Yes	\$3,000	The attackers demanded a ransom, which Krypton declined to pay. Krypton tried to turn to the PoS consensus mechanism to prevent future attacks, but the project was terminated a few months later.
Terracoin (TRC)	Jul 24, 2013	Yes	Unknown	Terracoin's price collapsed. The exchange Bter annouced that the attacker withdrew about 50 BTC value before the account was disabled.
Feathercoin (FTC)	Jun 8, 2013 - Jun 10, 2013	Yes	\$1,400	Feathercoin later adopted an Advanced Checkpointing (ACP) fea- ture to protect against 51% attacks. The checkpoint master node is deployed and maintained by the lead FTC developer.
Coiledcoin (CLC)	Jan 6, 2012	Yes	Unknown	The 51% attack killed CoiledCoin for non-financial reasons. Some community members accused Luke-Jr, a Bitcoin Core developer and the founder of Eligis mining pool, of using the pool resources to at- tack Coiledcoin. Luke-Jr denied it. But he stated that CoiledCoin was a scam that would discredit and harm Bitcoin's reputation.