DeFi Primitives Without Parallels: An Overview of AMMs, Consensus Algorithms, and Flash Loans

Anthony Hoang*

Renita Murimi**

Abstract

Decentralized finance (DeFi) has been hailed as a bold direction in the future of fintech, bringing with it the promise of democratized access to capital. The appeal of DeFi lies in its potential to rethink current financial paradigms and introduce new ones. The focus of this paper is on the elements that set DeFi apart from traditional elements of banking and finance. While there are many such elements, we sought to uncover the constructs that have little to no semblance with traditional financial processes. Specifically, we focused on automated market makers, consensus algorithms powering the blockchains, and flash loans that stand in stark contrast to order books, centralized governance, and collateral-based lending, respectively. This paper provides an overview of these unique DeFi primitives that have little in common with the world before blockchain, and present implications for responsible innovation in blockchain-powered fintech applications.

1. Introduction to DeFi

Decentralized Finance (DeFi) is a financial model that uses the distributed ledger architecture of blockchain to record transactions in a trustless, immutable, and decentralized environment. Compared to a traditional financial system, exemplified by the large centralized institutions such as banks and exchanges, DeFi works by removing institutional intermediaries that have been traditionally used to mitigate the information asymmetry and risk associated with financial transaction between multiple stakeholders. The removal of these intermediaries is facilitated by cryptographic trust, which is enabled due to algorithms that guarantee immutability, verifiability, and distributed access to records that are stored on the blockchain. Thus, DeFi has the potential to reduce transaction times as well as costs, thereby broadening financial inclusion, facilitating open access, and encouraging permissionless innovation. (Popescu, 2020).

Initially blockchain applications were driven significantly by the surging popularity of blockchain's first application - the Bitcoin cryptocurrency. Since 2008, when Satoshi Nakamoto made his Bitcoin whitepaper public on a cryptography mailing list, its underlying blockchain technology has seen immense growth, both in the field of cryptocurrencies and beyond. The net worth of cryptocurrencies exceeds more than a trillion dollars, and cryptocurrencies continue to feature as a prominent driver of DeFi applications. DeFi's architecture originates from decentralized applications (DApps) that are programmed to run on distributed networks.

*Anthony Hoang, Graduate Student, Gupta College of Business, University of Dallas; ahoang@udallas.edu.
**Renita Murimi, Associate Professor of Cybersecurity, Gupta College of Business, University of Dallas, and the Entropy Foundation, (972) 721-5058; rmurimi@udallas.edu.

Leveraging the open-source architecture and the use of cryptographic mechanisms on the blockchain, DApps enable the development of "token economies" for the creation and transfer of tokens, and possess the ability to adapt to varying contractual needs of the token stakeholders while still conforming to the underlying consensus protocols. Recent work in Chen and Bellavitis (2019) has created a taxonomy of DeFi business models which fall into four categories- decentralized currencies, decentralized payment services, decentralized fundraising, and decentralized contracting (Chen and Bellavitis, 2019).

Decentralized payment services offer the promise of financial inclusion to unbanked and underbanked people, while also expanding the scope of a cashless society. Decentralized payment services, such as Libra and the Bitcoin Lighting Network, offer the potential for low-cost, instant, and global payments. Decentralized payments services have addressed significant points of friction associated with the cost, time, and scale dimensions inherent in traditional payment services. Libra was first proposed by Facebook in 2019 as a global currency for payments and remittances. Later renamed the Diem Association, revisions to the initial Libra framework included support for tethered stablecoins and multi-currency coins, as well issues about monopoly, fraud, and compliance. Similarly, the Bitcoin Lightning Network offers a decentralized payment system with greater speed and anonymity than is possible with the underlying Bitcoin blockchain.

The token economy has been successfully leveraged for fundraising with blockchain counterparts. An Initial Coin Offering (ICO) offers institutional and individual investors a way to invest in digital tokens using the blockchain technology, much like IPOs are used to offer shares of a private corporation to the public with the goal of raising funds. Unlike IPOs, ICOs are characterized by the use of white papers, online sources (websites, blogs, social media) and code repositories such as GitHub to disseminate information about the technology and influence investors (Samieifar and Baur, 2020). ICOs create opportunities for newer cryptocurrencies or assets to be traded in the form of tokens in exchange for fiat currencies or existing, steady cryptocurrencies such as Bitcoin or Ethereum (Fenu et al, 2017). While the legal issues surrounding ICO regulation continue to be debated (Conley 2017), several ICO success stories exist with a few even making it to the rare ranks of unicorns, or private startups with a valuation of more than a billion dollars.

One of the drivers of DeFi adoption in various domains is its underlying layer of smart contracts. Smart contracts are software programs that automatically execute when pre-specified conditions in the protocols are met. Assets are represented as tokens, and each token's rights and responsibilities are translated into software code and placed on the blockchain. Assets could represent a wide range of entities, including but not limited to currency, real estate, business processes or intellectual property. The primary driver of smart contracts is thus the tokenizability of assets which represents the ability of an asset's features to be translated into software code for storing on a blockchain. Ethereum launched the capacity to create smart contracts, which ushered in the use blockchain for a wide range of applications in the digital asset economy.

In the rest of this paper, we will proceed to highlight multiple other unique features of DeFi in three broad categories. Section 3 describes automated market makers (AMMs) and decentralized exchanges (DEXs), along with an overview of the evolution of Uniswap, an AMM algorithm that powers the largest DEX for cryptocurrency exchange. Section 2 provides an overview of consensus algorithms, which form the foundation for block creation and ensuring that the entire community of nodes has information about the state of the blockchain. Consensus algorithms have evolved over the years, beginning with the first proof-of-work (PoW) algorithm used in Bitcoin. Section 4 discusses a novel variant of blockchain lending, exemplified in the form of flash loans and discusses the origins, processes, and attacks on flash loans. Section 5 presents implications for the future of DeFi by analyzing its potential for social good, while also offering caution regarding the development of secure smart contracts and DeFi protocols.

2. Algorithmic exchange mechanisms: AMMs and DEXs

AMMs are algorithmic traders that manage the liquidity in markets by setting the prices of assets. In the absence of a centralized entity to set these prices and maintain liquidity, a properly-designed AMM is able to withstand market fluctuations and manipulations. AMMs were initially analyzed for their application in prediction markets, where the outcomes of events are traded in exchanges in a binary manner (Othman et al, 2013). Prediction markets have found widespread application in diverse domains such as weather forecasting, political betting, and commodities trading. The most popular algorithm underlying AMMs for prediction markets is Hanson's Logarithmic Market Scoring Rule (LMSR) (Hanson, 2003). In a prediction market, accurate predictions about the outcomes of an event are the fuel that drives the prediction market engine. However, accurate predictions are hard to come by. Individual biases, information asymmetry, and misinformation all impinge upon the accuracy of predictions. Hanson's LMSR seeks to coax accurate predictions using financial incentives, by providing a logarithmic "score" to denote the accuracy of a prediction and to reward the prediction.

More recently, AMMs are being used in decentralized exchanges, where their algorithmic ability to continuously regulate the liquidity in the market with trading fees and incentives has been used in applications such as Uniswap, Curve, and Balancer. Unlike Hanson's LMSR that has been used predominantly in managing prediction markets, AMMs used in DEXs use constant function market makers. For example, Uniswap uses a function given by $x * y = k$, where x and y denote the quantities of two assets in the liquidity pool such that their product is a constant k. At each trade, the values of x and y change, thereby impacting their prices, which are all regulated through smart contracts on the Ethereum blockchain. An alternative to the constant product market maker functions is the constant mean market maker function which has been used in the Balancer DEX. Balancer differs from constant product AMMs like Uniswap, in that it allows for more than two assets to be maintained in a portfolio. Further, unlike traditional index funds where portfolio managers receive fees for maintaining the portfolio, the Balancer DEX requires traders to pay fees to the users to rebalance their portfolio. Assets are associated with weights and fees, and the Balancer AMM causes each trade to maintain a value in each asset of the portfolio that is equal to a constant. While Uniswap and Balancer feature liquidity pools that contain tokens of differently structured assets, their difference also leads to volatility in the proportion of assets and their prices causing phenomena such as slippage and impermanent losses. In contrast, the Curve DEX features liquidity pools of assets that are alike – stablecoins or wrapped bitcoin. Thus, the Curve market is less susceptible to impermanent loss due to fewer fluctuations in stablecoin prices, and can thus offer smaller fees to users.

DEXs are autonomous applications that exist within blockchains to allow users to trade without having to relinquish their funds to any intermediary party. Essentially DEXs allow users to mimic peer to peer trading by creating a peer to contract scenario where the trade is instantly accepted once the requirements of the contract are fulfilled. DEXs were initially created to eliminate the need for any intermediary to supervise and approve trades made within a transaction. The formation and maintenance of liquidity pools is enabled by staking, which can be in the form of liquidity mining or yield farming. Users get interest (yield farming) or rewards (liquidity mining) while the protocol gets liquidity to be able to function. Staking denotes the mechanism of giving up (locking up) assets to a specific pool within a protocol and farming/mining refer the rewards gained because of this transaction. Yield farming is a process that allows cryptocurrency holders to earn rewards on their holdings. An investor can deposit cryptocurrency into a lending protocol and earns interest from the trading fees within the protocol.

There are three types of DEXs that form a kind a spectrum where some protocols are completely decentralized while others are only quasi-decentralized. Traditional centralized exchanges such as banks use order books to log and keep track of every transaction that runs through the exchange. The first DEXs used on-chain order books which keeps the order books split throughout different network nodes that are assigned to maintain all transactions and requires miners to confirm each transaction. Examples of on-chain order book exchanges are Bitshares and StellarTerm. The next generation of DEXs use off-chain order books that recorded all transactions hosted in a centralized entity. A few examples of off-chain order books are Binance DEX and EtherDelta. These DEXs are closer to entities such as banks. The third kind of DEXs are automated market makers (AMM's). Automated market makers have no need for books or record keeping as they utilize smart contracts to form liquidity pools that automatically execute trades if certain parameters are met. AMM's are the backbone of true DEX's where a counterparty does not exist. Instead, users execute a trade against the liquidity in the liquidity pool that is managed by an algorithm that manages the pool. Liquidity pools are key components of DEXs, and are broadly described as the collections of funds locked in a smart contract. They are used to facilitate decentralized trading, lending, and many other functions. For example, the liquidity pools for the DEX PancakeSwap are called Syrup pools. Syrup pools uses PancakeSwap's token, CAKE, to allow users to stake their coins. These coins can earn "syrup", a portion of the block reward, by being randomly selected by the PancakeSwap team, or distributed to users that vote for different community projects to maintain a sort of user governance.

One of the metrics used to characterize DEXs is TVL, or Total Value Locked, which represents the number of assets being staked in a specific protocol. This value is not meant to represent the number of outstanding loans, but rather the total amount of underlying supply that is secured by the DeFi application. TVL is best utilized as a metric for overall health of a DEX and and is used to determine if a DeFi asset is overvalued or undervalued. In theory, if a TVL ratio is under 1, it is probably undervalued. For example, Uniswap has a market cap of $11 billion and a TVL of $7.6 billion. The TVL ratio of Uniswap is 1.45 and shows that Uniswap is currently overvalued. On the other hand, PancakeSwap has a market cap of $2.9 billion and a TVL of $6.5 billion. The TVL ratio of PancakeSwap is 0.45 and shows that PancakeSwap in currently undervalued.

The Evolution of Uniswap

AMMs for DEXs are a relatively new market instrument for cryptoasset exchanges. As the DeFi market matures, the range of AMMs used to regulate these markets is expected to vary in terms of numbers of assets, trading fees, volatility, and incentives for participation. DEXs are also subject to unique constraints posed by the nature of AMMs, such as impermanent loss and slippage. When users provide liquidity to a pool and they receive lesser value at withdrawal of their liquidity, it is called impermanent loss. Some solutions to counteract impermanent loss include the requirement of trading fees, use of bounded ranges for liquidity, or the use of stablecoins and wrapped assets. On the other hand, slippage refers to the difference between the quoted price and the executed price of a token on a DEX. Large orders tend to create low liquidity, which in turn, induce slippage in the exchange. Some solutions to counter slippage include splitting larger trades into smaller chunks, choosing slippage tolerance at the time of the trade, increasing the gas associated with a transaction, and the use of blockchain architectures that do not clog the Ethereum blockchain.

First launched in 2018, Uniswap is the largest DEX on the Ethereum blockchain. Uniswap allows for on-chain token exchanges in pairs, where the quantities of these tokens and their corresponding prices are dictated by the constant product AMM algorithm. In this section, we focus on the Uniswap AMM, which is the largest DEX on the Ethereum blockchain. Recent literature has analyzed the Uniswap DEX in terms of its alignment to true market price (Angeris et al, 2019), liquidity provision (Neuder et al, 2021), losses and risk profiles (Aigner and Dhaliwal, 2021). Below, we provide an overview of the evolution of the core algorithm used in the AMM for Uniswap.

Uniswap v1: Launched in 2018, the first version of Uniswap allowed for token exchanges between any ERC-20 token and ETH. Thus, Uniswap allowed for multiple ERC-20/ETH exchanges, and users could contribute to the liquidity pools (reserves) of any of these exchanges. As the tokens in an exchange were being sold/bought, their share in the liquidity pool was altered, which in turn, was reflected in the price of the tokens. An exchange between two tokens A and B would have to proceed in two stages, $A \rightarrow ETH$ and then, $ETH \rightarrow B,$ where ETH acted as the “bridge currency”. While there were no fees for listing tokens or using the Uniswap platform, a small liquidity provider fee was taken out of each transaction and added to the reserve. This ensured that the reserve size kept increasing with every transaction, even though the ratio of tokens in each exchange depended on the size of the trades being executed. The process of the two-stage exchange imposes varying degrees of impermanent losses, depending on the correlation between the two types of tokens. Thus, the exchange of two stablecoins would be subject to smaller impermanent losses, compared to the exchange of tokens that are more closely related to ETH. Further, the two-step exchange process introduces twice the trader fees, that results in twice the amount of slippage.

Uniswap v2: While continuing to use the same constant-product AMM formula as in Uniswap v1, the second version of Uniswap's algorithm differed significantly from Uniswap v1 in that it no longer required ETH as a bridge currency in exchanges between two ERC-20 tokens. Launched in 2020, another significant difference in the implementation was initiated with respect to arbitrage-induced manipulation that was possible with Uniswap v1. Angeris (2019) suggested that since the ratio of tokens in an exchange was correlated with price, Uniswap acts as "approximate price oracle". This also made it possible to manipulate the prices offered by Uniswap v1 by buying a token from the exchange, settling it with inflation, and selling back the token to the contract. Uniswap v2 made it harder to be manipulated in this manner, by incorporating a weighted average of the prices. Specifically, instead of using a single price recorded by the contract prior to a transaction, Uniswap v2 kept track of the cumulative sum of prices prior to a transaction and this sum was weighted by the time elapsed since the last time instant the price was updated in a block. This mechanism of using the difference between two accumulated prices (p_1, p_2) between two time-instants t_1 and t_2, which was then divided by the time elapsed $(t_2 - t_1)$ is called the Time Weighted Average Price (TWAP). The TWAP model was harder to attack by way of manipulation, but at the same time created a price oracle that was less up-to-date, since it relied on a weighted average of accumulated price differences, and not the instantaneous price. Other changes over Uniswap v1 include a change of programming language (v1: Vyper, v2: Solidity), the ability to swap tokens within the same transaction without having to pay to buy a token first, and the requirement for wrapped ETH tokens.

Uniswap v3: The current version of the Uniswap v3 price oracle algorithm was launched in 2021 and introduced further changes to the AMM algorithm. While v2 used accumulated prices, it did so externally. Uniswap v3 accumulated the prices on-chain and offers the ability for external contracts to incorporate these on-chain calculations. Further, while v2 used an arithmetic mean of accumulated prices, v3 introduced the geometric mean TWAP. Multiple reasons were cited for the use of the geometric TWAP, including a closer alignment to the market price, ease of implementation, efficiency of storage (since the geometric mean uses logarithm of the price, and not the actual price). Another significant difference was the use of concentrated liquidity that offered users a range of prices within which they can bound the liquidity, instead of offering the entire liquidity pool $(0, \infty)$ in Uniswap v1 or v2.

In addition to Uniswap, other examples of DEX protocols on BSC that allow for the decentralized exchange of tokens are Pancakeswap, BurgerSwap, and BakerySwap. Each provide their own set of rewards to differentiate them. For example, Pancake Swap is a DEX that allows users to trade cryptocurrencies and tokens on the Binance Smart Chain. Users connect their wallet applications to Pancake, and are able to swap crypto tokens on the Pancake DEX using smart contracts. In addition to offering its own native token called CAKE, Pancake Swap offers gamified experiences that incentivizes users to stake their coins to try to earn a profit. For example, PancakeSwap allows users to wager on whether the price of the Binance token (BNB) will increase or decrease in a certain window of time. A summary of well-known DEX swap exchanges is in Table I.

3. Unity within and outside the blockchain: consensus algorithms and oracles

This section provides an overview of a fundamental building block of blockchain – the consensus algorithms. Consensus algorithms are essential to keeping track of the evolving blockchain, as new transactions and new blocks are processed. Keeping true to the decentralized nature of blockchain, consensus algorithms leverage the power of the entire network to agree on the state of the network on the distributed ledger. Further, this section introduces the notion of blockchain oracles, which are critical to building a link between the trusted blockchain and the untrusted external environment. Oracles fetch information from the external world and supply it to the blockchain for various applications. Together, consensus algorithms on the blockchain and oracles that bridge the on-chain and off-chain environments are essential to the design and development of robust DeFi protocols.

3.1 Oracles

Blockchain oracles function serve as an intermediary between the trusted blockchain environment and the untrusted data sources outside the blockchain (Murimi and Wang, 2020). While cryptographic mechanisms enable trust-free operation on the blockchain, for complex environments such as DEXs there is a need to incorporate information that is external to the blockchain. Blockchain oracles help to bridge this gap between the blockchain and data sources around it. The data from external sources is fetched by oracles, and is then used to activate smart contracts that link a diverse set of stakeholders in the blockchain environment (Wohrer & Zdun, 2018). For example, in a DEX, arbitrageurs interested in trading opportunities due to price discrepancies of tokens in liquidity pools in different markets are dependent on data from external websites (untrusted environment). Accuracy of current market data, thus, is a prime factor, in the success of the arbitrage activities. As a trusted entity, the blockchain oracle obtains this information and supplies it to the nodes on the chain, and serves as a crucial link between the two environments.

While the above example describes an oracle that fetches commodity prices, other types of oracles exist. For example, in a hedging model, nodes in a blockchain (trusted environment) might be dependent on weather data from an external website (untrusted environment) in order to predict future prices for an agricultural commodity, while other oracles might be responsible for tracking polling and prediction markets (Al-Breiki et al, 2020). In this paper, we limit our discussion to price oracles that are involved in supplying price feeds for liquidity pools on DEXs.

Thus, a price oracle may be defined as any tool that is used to view the price of an asset, in this case, the price of different tokens. Price oracles gather data from DeFi protocols to determine a price based on the oracle’s algorithm. Price oracles can be categorized based on how they source their data. Centralized oracles process the data through a single source and decentralized oracles rely on multiple sources. While decentralized oracles use a variety of sources to compile data, centralized oracles have the risk of collapsing the DEX because an attack on the centralized oracle would negatively impact every smart contract that relies on the price oracle.

3.2 Consensus algorithms

Proof-based consensus algorithms are called thus, since a node in this category has to compete with other nodes and prove it is more qualified to commit transactions. The consensus layer forms the core of the consensus process used to determine the validity of the block data by the highly decentralized nodes. The main consensus mechanisms are Proof of Work (PoW), Proof of Stake (PoS), and Proof of Authority. Variations of these basic consensus mechanisms include Delegated Proof of Stake (DPoS), Proof of Staked Authority (PoSA), and leased PoS, among others. Other proof-based algorithms that use time, space, and reputation as a means of achieving consensus include proof of space (also known as proof of capacity), proof of elapsed time, proof of activity, proof of burn, and proof of importance. One of the primary goals of achieving consensus is for the entire network to believe the same information about the state of the network. This is related to the 51% attack, where one or more miners gain control over more than half of the network's mining power. This has significant implications for the state of the blockchain, since the nodes that control the majority of mining power can determine which transactions to add to the network, order of transactions, block other miners, and enable double-spending. Double-spending refers to a network attack when the currency is being spent more than once, and occurs due to inconsistent information about the state of the network.

Below, we provide a brief overview of these consensus algorithms, and feature the applications that leverage them.

3.1 Proof of Work (PoW)

The PoW algorithm, implemented in Bitcoin, is based on the premise of winner-takes-it-all. Miners compete against each other to generate a hash that matches a current "target". This hash is a string of characters with a pre-determined set of zeros leading the string of characters. The mining process is incentivized by rewarding the first miner to obtain this target hash with Bitcoin. Once a hash is generated – which takes about ten minutes – a new target is generated for which miners compete by spending their computational power all over again. This expenditure of computational power is the "work" that miners exert, and so the nodes prove that they completed the work necessary to earn Bitcoin. Although popularized by use in Bitcoin, Bitcoin Cash, Litecoin and Monero, PoW has its origins in applications that were built to deter email spam (Back, 2002).

3.2 Proof of Stake (PoS)

In the PoS consensus mechanism, nodes that want to validate a transaction for adding it to a block are chosen in accordance with the amount of ether (ETH) that they stake on the blockchain. These validator nodes do not know if they are going to be picked to approve a transaction; instead, they are picked randomly. Nodes that approve a transaction erroneously stand to lose a portion or the entire amount of their stake, which is a minimum of 32 ETH for the Ethereum blockchain. The staked ether of all validators determines the rewards that validators receive. Further, as other DeFi protocols such as Maker and Aave use ETH as collateral, the total amount of ETH available for purchase on the network reduces, increasing the price of ETH further.

3.3 Delegated Proof of Stake (DPoS)

In a variation of PoS, delegated PoS leverages an additional role of "witness nodes". Nodes vote to choose witnesses according to their stake; thus, more stake translates to more votes that are available for a node. The witnesses are chosen to validate transactions and a smaller subset of these witness nodes are incentivized for their work. The election for witness nodes is a continuous process, since witness nodes who behave maliciously or who fail to produce a node are required to relinquish their role as a witness node. The smaller set of witness nodes that validates transactions ensures a faster consensus process, enabling higher transactions per second than PoS or PoW.

3.4 Leased Proof of Stake (LPoS)

Since mining is computationally expensive either in terms of electricity or stake in the blockchain, leased PoS offers an alternative where users lease their stake to miners in return for a portion of the mining rewards. Like its parent PoS algorithm, LPoS validating nodes are still chosen from the network based on their stake. While the leasing process institutes an additional type of transaction (leasing), the number of tokens in LPoS is fixed since the mining does not introduce additional tokens to the token pool since existing tokens are leased and released when the mining is completed. Also, since fewer nodes are involved in the validation process, it offers the benefits of speed and lower computational processing in the network.

3.5 Proof of Activity

In an extension of PoW, PoA employs additional vetting procedures to confer mining rights on nodes. PoA is a combination of PoS and PoW, where mining proceeds in two phases. In the first phase, miners compete to be the first to solve a puzzle and claim their reward. In this first preliminary phase, the mined blocks are almost blank. They are simply templates with header information and the mining reward address, and do not contain any transactions. Thus, the first phase is similar to the PoW protocol. In the second phase, the PoS protocol is leveraged where the header information from the mined pseudo-blank block is used to select a random group of validators to sign the block. This group of validators are chosen according to the stake they have placed in the system. When the entire set of chosen validators signs a block, it gets added to the blockchain. The signing is time-sensitive; blocks that are not signed by all validators are discarded in favor of the next winning block. A new set of validators is chosen to sign this new winning block. PoA offers rewards in the form of network fees to the winning miner and the set of validators who signed the block.

3.6 Proof of Authority (PoA)

Unlike PoW where nodes compete based on mining power or PoS where nodes compete based on their stake in the network, PoA is based on the concept of staking reputation. The identity of a node is associated with a reputation, and validating nodes are randomly chosen based on their reputation. US notaries in good standing form the group of validators, after being subject to identity verification. Validators are further chosen based on criteria such as their knowledge of blockchain, degree of commitment, merit and interest in being a validator.

Candidates interested in becoming a validator are evaluated by their performance on the PoA public forum, participation on the testnet, and ultimately an endorsement from a current validator for inclusion on the ballot.

3.7 Proof of Reputation

A variation of PoA, called Proof of Reputation (PoR), stakes the reputation of companies instead of individuals. Since the reputation of a company is associated with the reputation of its employees and market cap, the verification process to choose a validating node is designed to increase the stake that validators place in the network. Validators, called authorized signers, have their information available on the blockchain. As such, PoR has been primarily used in private blockchains where the participating entities are only allowed to participate in the blockchain after an initial vetting process.

3.8 Proof of Staked Authority

While PoA relies on a limited set of validators to achieve efficiency in processing speed, it also results in a form of centralization by limiting the responsibilities of validation and chain creation to a limited set of nods. To avoid the potential pitfalls of fraud targeted at or generated from a centralized architecture, PoSA uses a combination of PoA and dPoS. The PoSA, used in the BSC, achieves this combination by still using a limited set of validators to produce blocks, but the choice of validators is based on the stake of each validator. The top 21 most staked nodes are chosen to be in the current validator set, and elections are conducted every 24 hours. Validators are incentivized for their work with transaction fees (generated from the transactions in a block) that are rewarded in the form of BSC tokens called BNB. These tokens can be used as fees to deploy smart contracts on the BSC, and are compatible for cross-chain communication between BSC and BC.

3.9 Proof of Importance

In PoI, several metrics such as the stake of individual nodes, duration of ownership of the stake, number and size of transactions in recent history, and ratings assigned to individual nodes are used to calculate the importance of a node. Individual nodes are called block harvesters, and the calculated value of the importance metric is used to select the node that will mine the next block. If the chosen block is offline, the task of mining the block is delegated to a different node. PoI avoids expensive computational procedures throughout the entire network that are inherent in PoW, and therefore uses less energy than PoI.

3.10 Proof of Space (PoSpace)

Unlike PoW where provers claim to find the target hash in order to earn cryptocurrency, PoSpace requires the prover to demonstrate that it has a required amount of storage space reserved. The motivation for PoSpace lies in non-trivial amounts of free disk space that most users have. Proof of Secure Erasure (PoSE), a related algorithm, relies on a prover to provide evidence that she has erased its memory of size N. These space-related proof algorithms are computationally more efficient than PoW, requiring less energy and avoiding the use of special-purpose mining equipment for users.

3.11 Proof of Elapsed Time (PoET)

Taking a different route to address the high energy consumption of blockchains running PoW and avoiding the staking processes of PoS, PoET is designed to reduce energy consumption by randomizing the choice of nodes that commit a block to the blockchain. Nodes are assigned random numbers that denote the amount of time that the node will sleep; the node with the smallest random number is one with the shortest sleep time and upon waking up, it commits a new block to the blockchain. Designed by Intel and used in private blockchain networks, the PoET algorithm is used in the Hyperledger Sawtooth implementation. Since nodes are assigned sleep time durations randomly, every node is offered a fair chance at being given mining rights without the need to stake existing currency or reputation (like in the PoS algorithm).

3.12 Proof of Burn

The transfer of cryptocurrency requires a verifiable sending address and receiving address without which the transfer might result in the currency being deposited to the wrong address or being deposited to a verifiably unspendable address, also called an “eater address”. While the former scenario has limited support for retrieving the funds, the latter results in a complete, irrevocable loss of funds which is also called “burn” of cryptocurrency. PoB is still experimental, and requires miners to burn some cryptocurrency to gain mining rights. The higher the number of coins indicates higher mining power, and higher potential rewards from the mining rights.

The next section highlights a unique DeFi application – flash loans, and compares them to traditional fiat loans, while also highlighting and analyzing several high-profile flash loan attacks.

4. A novel DeFi application: Flash loans

Historically, loans were viewed through the lens of social, religious, and political perspectives about the kind of profits that a lender could make through the mechanism of loans. While certain societies were altruistic about the interest rate, collateral, and duration of loans, others imposed varying rules on loans made to individuals within and outside the community, and in general, outsiders were treated with different, stringent rules concerning the lending process (Rajan, 1992). Weber’s last theory of capitalism hypothesized the state as a rational/legal entity that provided a standardized method of investments, banking, taxation, currencies, among others thus eliminating adhoc processes concerning loans (Collins, 1980). The progress made in DeFi, however, represents a different shift back to decentralized processes. It remains to see if decentralized systems of banking and finance complement or diversify the role of centuries of state-backed financial, social, and political interventions. In this section, we examine the features of traditional loans and compare them with the features of flash loans, and highlight challenges and opportunities for growth in flash loan instruments.

4.1 Traditional loans

Collateral and loans have traditionally complemented each other in the quest to reduce market risk. The market for loans is characterized by a fundamental information asymmetry, and the use of mechanisms such as collateral, loan interest, and credit scores serves to partially mitigate the risks posed to the lender by this asymmetry (Chen et al, 2021). A detailed study of the various factors involved in the use of collateral was performed in Jimenez at al. (2006). Specifically, they studied the role of credit quality of the borrower, type of lender specialization, the competition in the credit market, duration of the loan, size of the loan, and the macroeconomic conditions. In particular, they examined two kinds of borrower-lender relationships: relational and transactional. Relational lending is characterized by a higher number of relationships and a greater concentration of those relationships in a reduced number of banks, while transactional lending is characterized by fewer relationships with each of many banks. In relational lending, the probability of the use of collateral decreases with greater existing trust due a longer borrower-lender relationship. The authors found that for long-term loans, relationship lending yields better loan terms than transactional lending. On the contrary, for short-term loans, they found that transactional lending was preferable compared to long-term loans. These findings build upon the work of Boot and Thakor (1994) which showed that collateral is a mechanism for reducing moral hazard, and Jimenez et al (2006) showed that the use of collateral was also a signal for decreasing adverse selection in markets.

While formal credit markets depend heavily on the use of collateral as a corequisite in the lending process, individuals in low-income communities that are heavily resource-constrained look to informal credit markets such as local money lenders where the interest rates are higher than in the formal credit markets (Caskey, 1997). A different lending economy has emerged in such regions in the form of microcredit or microloans, where loans are made to low-income women in rural communities individually or in small groups. These microloans are characterized by the following attributes: short term duration, frequent interest payments, and the use of social collateral instead of physical collateral (Cheung and Sundaresan, 2006). Further work in Elahi and Rahman (2006) classified microcredit into several categories, and attributes the above definition of a microloan to the category of Grameen-credit. Grameen-credit was first designed and popularized by Muhammad Yunus, founder of the Grameen bank who implemented microcredit loans in an effort to reduce the barriers for rural, poor women to access credit.

Earlier work in Chan and Kanatas (1985) theorized that an environment with risk-neutral transactions and complete information symmetry would make the use of collateral obsolete. The use of collateral by lending institutions, specifically banks, has been examined in Blazy and Weill (2013). Here, the authors describe collateral as a mechanism to reduce the problems caused by adverse selection before the loan is made, and the problems caused by moral hazard after the loan is made. Work in Berger et al (2016) found the higher the liquidity of the collateral, the lower its association with the risk and therefore performed better than loans with illiquid collateral or no collateral. A distinction between the cost and the value of collateral from a borrower’s perspective is investigated in Ninimaki (2011), where the author showed that two factors - market conditions and the probability of success of a project for which a borrower is applying for a loan – are significant in determining the cost of collateral to the borrower.

Figure 1 represents a schema of loans classified according to their collateral types. Tangible collateral is exemplified in the form of discrete assets including but not limited to real estate, cash, and inventory. On the other hand, intangible collateral can be classified into two further subcategories. The first of these is comprised of non-discrete assets such as intellectual property, domain names, and customer leads. The latter subcategory belongs to the mechanisms for lending in the crypto-economy, such as flash loans. For a detailed discussion of intangible collateral types, the reader is referred to Loumioti (2012). Flash loans, in essence, are intangible due to the nature of the crypto-assets being borrowed/lent and cannot be grouped into the categories of either discrete or non-discrete assets due to their tendency to be grouped as alternative assets (Bartolucci and Kirilenko, 2020).

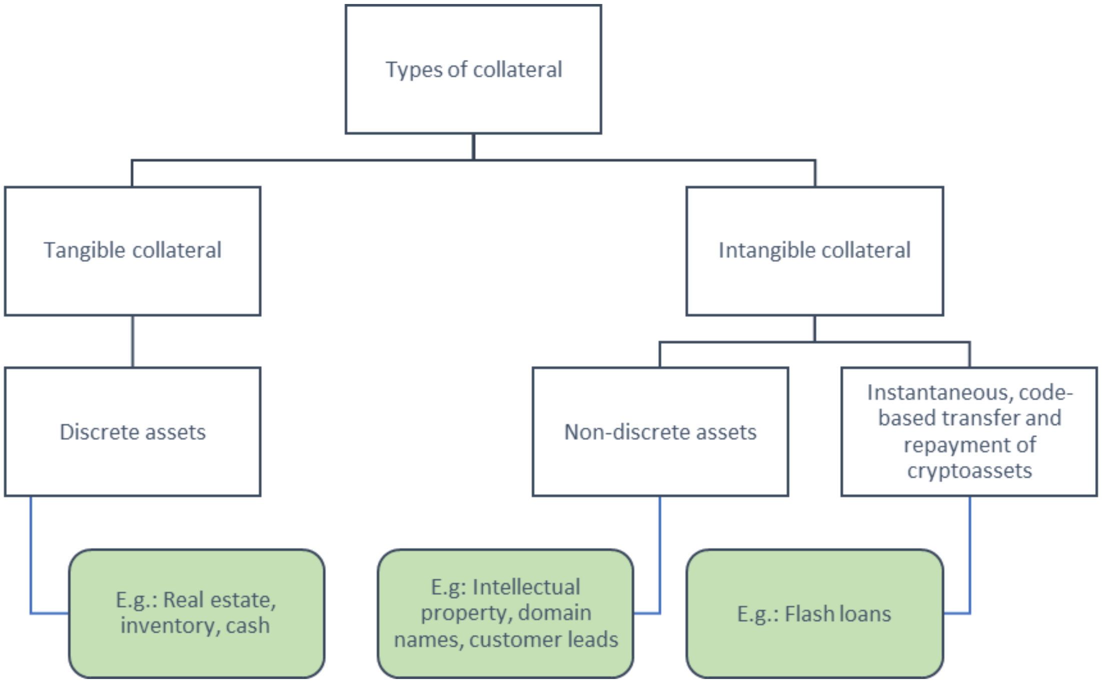
Types of collateral
Tangible collateral
Intangible collateral
Discrete assets
Non-discrete assets
Instantaneous, code-based transfer and repayment of cryptoassets
E.g.: Real estate, inventory, cash
E.g: Intellectual property, domain names, customer leads
E.g.: Flash loans

Figure 1. A schema of collateral types in flash loans.

4.2 Flash loans

Flash loans are a popular construct in DeFi protocols that support arbitrage opportunities. Flash loans are uncollateralized loans, i.e. they do not require any existing collateral or assets for users to borrow as much as they want. The caveat is that the loan must be repaid within the same transaction, the duration of which is typically a few seconds. This allows arbitrage traders to loan cryptocurrency from one market, and sell it to another market for profit, and repay the loan in full within the same transaction. An example of a flash loan is described in Wang et al (2021), where they reported an arbitrage opportunity that spanned multiple DEXs (dYdX, Balancer, Uniswap) and the conversion process between currencies during a single transaction in January 2021 was as follows: (1) Borrowed 1.13 ETH from dYdX. (2) This amount was converted to 345 (Livepeer) LPT tokens in Balancer. (3) These 345 LPT tokens were traded in a Uniswap liquidity pool, netting 1.46 ETH for the trader. (4) Finally, the trader returned 1.13 ETH to dYdX. This arbitrage profited the trader an amount of 0.33 ETH (1.46 ETH – 1.13 ETH) (around 538 USD at the time) with a gas fee of 0.05 ETH.

Flash loans, although uncollateralized, address the issue of information asymmetry using the reversal feature of the smart contract. The premise of a flash loan is that the loan has to be returned within the same transaction, otherwise the loan is reversed at the end of the transaction. This reversal, enabled through the code, offers parallels to that of conventional collateral. Conventional collateral-backed loans assume that ownership of the collateral may be transferred to the lender, either partially or wholly, in the event of non-repayment of the loan. Further, in the event of non-repayment, if the collateral in a conventional loan is much smaller than the loan amount itself, the lender's risk profile is significantly increased. In contrast, in a flash loan, despite the lack of collateral, the lender is guaranteed to recuperate the funds at the end of a transaction either through repayment or non-repayment on the part of the borrower. Thus, the lender's risk profile is significantly reduced in a flash loan, due to the automatic reversal of the funds enabled by the smart contract.

Despite the promise of flash loans to offer easy access to capital with low risk, the DeFi market has recently been faced with a barrage of flash loan attacks. Since flash loans are extremely sensitive to the market, these attacks are not simply arbitrage opportunities. Rather, they are purposeful manipulations of the market resulting from exploits to code, hacks, or security breaches that create a destabilization of the market. For example, a flash loan attack on PancakeBunny, a yield aggregator of PancakeSwap, crashed the market, by devaluing a coin, originally valued at $145, became valued at $20. Alongside 8 consecutive flash loans, $45 million dollars were taken and none of it was recovered after the fact (Crawley, 2021).

4.3 Flash Loan Attacks

Flash loan attacks are essentially, attacks on price oracles which are the external entities that use data feeds to determine the price of tokens. These price oracle attacks benefit from any forceful destabilization of the market and this destabilization comes from the systematic risk in centralized oracles. In this section, we analyze recent flash loan attacks on the Binance Smart Chain (BSC) and provide a categorical overview of the different kinds of flash loan attacks.

The Binance Smart Chain, or BSC, was created by the developers of Binance Chain as a platform for cryptocurrency exchange by using smart contracts and virtual machines. Since BSC is designed to run in parallel to the Binance Chain, users are able to transfer their assets seamlessly between the two based on their preference of complexity. The difference comes through BSC's smart contract functionality and the consensus algorithm it uses. Key differences between the Binance chain and the BSC include the type of consensus algorithm (dPoS versus PoSA in BSC), number of validator nodes (11 versus up to 21 in BSC), and mean block processing times (1 second versus 5 seconds in BSC). Further, BSC offers the ability to facilitate dual-chain communication and supports high-performance dApps compared to the Binance chain that offered limited cross-chain support.

We examined 11 different flash loan attacks that were targeted against the BSC over the period of a year. Our evaluation of these flash loan attacks on the BSC reveals that flash loan attacks fall into three major categories (see Figure 2). These categories are bugs in the smart contracts, token manipulation, and faulty price oracle design. The first category covering smart contract bugs represents software bugs in the code, resulting in security flaws or protocol implementation flaws. The second category surveying token manipulation comprises of any artificial control over tokens that give the attacker an advantage by inflating the market. This category can be further divided into two sub-categories of minting-based or burning-based token manipulation. Minting is the creation of tokens while burning is the deletion of tokens. The third category of flash loan attacks is faulty price oracle designs that cause mis-valuation of the tokens. The rest of this section details the prominent flash loan attacks in each of these categories.

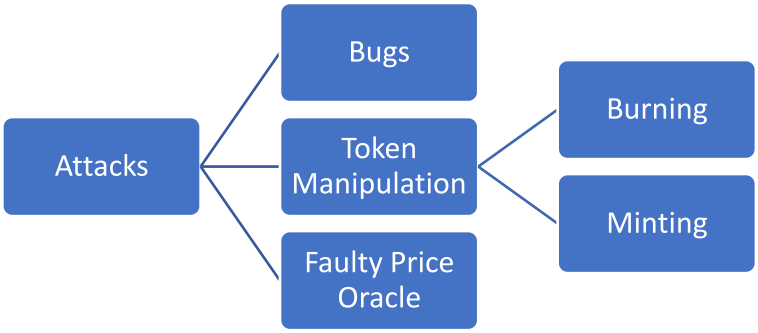

Figure 2. Categories of flash loan attacks on the BSC

Bugs in the smart contracts

Belt Finance: Belt Finance uses liquidity pools that are linked together in terms of price. This attack exploited a bug within this linked interaction by depositing in one pool and withdrawing in another creating a large discrepancy in price and swapped the Binance stable coin (BUSD) to stable coin USDT. The resulting fallout led to a withdrawal of BUSD due to the price discrepancy, which was followed up with seven flash loans profiting the attacker to the tune of $6.3M (Bitcoin News and Reports, 2021).

bEarn Finance: This attack exploited a bug in the withdraw function allowing a larger withdrawal than normal. The attacker followed up with 30 consecutive flash loans and repaid them taking advantage of this rare “arbitrage” position with a profit of $10.8M (Young and Baird, 2021).

Poly: The attacker was able to exploit a bug in the Poly network system resulting in the attacker gaining elevated internal permissions that gave the hacker permission that was only allowed internally. With these elevated privileges, the attacker was able to obtain a key that they used to call a function called “EthCrossChainManager” which transferred the ownership of smart contracts to the attacker. The attacker earned a total of $600M but promised to return the money with a warning to the DeFi community (Browne, 2021).

Token manipulation by burning

Pancake Bunny: The attacker used the function “getReward()” to mint BUNNY tokens and later burned them. This caused a major disturbance in the market causing the market value of

BUNNY to crash from $145 to $20. The attacker followed up with 8 consecutive flash loans and paid back the loans with a profit of $7.2 million. (Crawley, 2021)

Spartan Protocol: This attack exploited a flaw in the liquidity share calculation in the protocol that allowed the attacker to execute the function that burned tokens. The attacker borrowed $61M in BNB and burned the tokens to create an imbalance in the market. The attacker then swapped tokens and repaid the loan making a profit of $30M (Bourgi, 2021).

Token manipulation by minting

bZx: This attack started off with a 10000 ETH flash loan from dYdX, which was used to then send money to other exchanges. The attacker then minted tokens to get more funds and then sent 1300 ETH to bZx destabilizing the ETH/WBTC pool to a 1:5 ratio. The attacker was able to convert 112 WBTC (earned by minting tokens) to 6871 ETH on Uniswap due to the artificially skewed prices and the attacker used that to pay back the initial flash loan (Heasman, 2020).

Impossible Finance: The attacker produced a fake token within the market allowing the attacker to trade for IF’s native token. The attacker then followed up with multiple flash loans, borrowing IF’s native coin and trading it for BUSD and then traded BUSD for BNB. The attacker minted a fake coin to create a destabilization in the market which drained the liquidity in the pool. The attacker came out with a profit of $500K (Young, 2021).

Burger Swap: The attacker made a fake BEP-20 token that was paired with the BURGER token. The attacker then minted fake tokens to create a rift between that trading pair. The attacker followed up with multiple flash loans to take advantage of this rift and paid back the loans with a profit of $7.2 million (Redman, 2021).

Bogged Finance: The attacker exploited a bug in the contract that allowed them to mint new tokens. These new tokens created inflation in the market, causing the market price go from $1.8 to almost $0. The attacker followed up with flash loans and paid them back with a profit of $3 million (Kahraman, 2021).

Faulty price oracles

Rari Capital: The attacker used a price manipulation attack to trick Rari Capital’s smart contract to misvalue the Interest Bearing ETH (ibETH) token and allowing the attacker to buy the tokens at a lower price. Compounded with a subsequent flash loan, the attacker traded the token and repaid the loans with a profit of $10M (Thurman, 2021).

ApeRocket: The attacker started a loan in AAVE and PancakeSwap and deposited tokens in their respective pools within ApeRocket. The attacker was able to borrow a large number of tokens which tricked the smart contract into minting SPACE tokens as a reward for the deposit. The attacker then swapped the rewarded SPACE token for CAKE and repaid the loan with a profit of $1.26M (Vermaak, 2021).

4.2 Discussion

While flash loans represent a revolution in uncollateralized lending with the maximum loan size dictated only by the size of the liquidity pool, the smart contracts that enable their execution render them vulnerable to attacks and exploits of the underlying protocols. Work in

Qin et al (2021) systematically analyzed two flash loan attacks, and provided a detailed timeline of the various arbitrage transactions that were set into motion during the course of the attacks. Further, the authors analyzed the opportunity loss suffered by the attackers, who could have reaped higher profits from their attack by leveraging efficient attack parameters. Their findings show that the foundational premise of DeFi which allows distributed innovation needs a stronger security posture for the underlying protocols, without which DeFi attacks such as those geared toward flash loans can undermine the crypto-economy. Further, existing regulations have to be rethought regarding their application to flash loans.

Since a flash loan is executed within a single transaction, and the lender effectively does not part with the funds, flash loans contradict the basic principles of taxation of interest (Rotfleisch, 2021). As highlighted in (Chohan, 2021; Caldarelli and Ellul, 2021), DeFi in general does not espouse the widely established Know Your Customer (KYC) and Anti Money Laundering (AML) frameworks that are used extensively in the fiat currency economy. Among the many aspects confounding the development of holistic regulation for DeFi, are the questions surrounding regulation of particular groups of DeFi entities. As noted in Salami (2020), regulation applicable to the developers of DeFi applications and users of these applications will be different from the regulation developed for banks and their customers. Further, the author illustrates a key existential point about blockchain technologies – since they cannot be shut down, how effective will regulation will be in any case? Still, solutions in the form of secondary liability that were first developed to address copyright issues on the Internet show promise for regulation of the online platforms that facilitate DeFi (Wright, 2020). Other novel approaches for regulation in DeFi include substituted compliance or equivalence where DeFi activities licensed in one jurisdiction reduce the supervision for that project in other jurisdictions, regulatory cooperation, and embedded regulation (Zetzsche, 2020).

5. Implications for the future of DeFi innovation

Having surveyed several DeFi primitives, in this section we provide an overview of opportunities for DeFi innovation in four categories: social good, systemic development, security-centered design.

a. DeFi for social good: Laboure and Braunstein (2021) opined that the promise of financial inclusion, greater transparency, and diversified financial services have the potential for reducing economic inequality. Although DeFi prides itself on decentralization offered by the underlying blockchain, the enormous crypto-ecosystem that blockchain fuels might benefit from a mix of centralized and decentralized regulation. The technical, social, and policy aspects of DeFi converge in widely varying applications that cater to different industry domains. The lessons learned from more than two centuries of centralized efforts and the growing decentralization efforts in IoT, computing, and blockchain offer insights that point to a fluid combination of centralized and decentralized processes for wider impact beyond a few crypto-native sectors.
b. Systemic thinking for the crypto ecosystem: In Werhane (2002), the authors argue for the need for systemic thinking using the example of the failure of Western lending practices in Bangladesh. When banks in Bangladesh moved to lending processes based on

collateral, that effectively shut out the vast majority of their population who did not have access to resources that could be put up as collateral. Another example of systemic thinking can be found in decentralized governance models, where local modes of governance are better than regional governance in areas such as allocation of local resources. Flash loans be another manifestation of decentralization in cryptocurrencies, where the collateral, penalties, and loan recuperation processes are managed by the smart contract itself.

c. A chance for reconceptualizing security: The rise of blockchain has been widely heralded as the Web 3.0. The original Web was designed primarily as a mode of communication, and security was mostly an afterthought, which is manifested in the patchwork of protocols that are being used to provide security for the various services and applications hosted on the Internet. The Web 3.0 gives us an opportunity to rethink security that is built inside the applications, and not just delivered as an add-on. Cryptographic primitives enabling the blockchain are one way that applications such as DeFi can promise secure transactions, however, that is only the start. The protocols connecting the various applications to the blockchains are at the forefront of the discourse about secure DeFi applications, and designing DeFi with inbuilt security will further enhance the relevance and applicability of newer services such as flash loans.

While smart contracts offer the inherent benefits of decentralization and immutability due to their foundations in the blockchain architecture, the security of smart contracts is only as good as that of the software programming languages that are used to write the code. Common constraints in the development of smart contracts include the choice of a Turing-complete or Non-Turing-complete language, development of Byzantine fault-tolerant algorithms to prevent intentional tampering, and the lack of anonymity due to the smart contract's deployment on a blockchain. Among other features related to memory usage and computability of functions, Turing-completeness refers to the ability of a computer program to not halt on its own (e.g. recursion functions). The Bitcoin scripting language is Turing-incomplete, while the Solidity programming language used in Ethereum is Turing-incomplete. The ability of a programming language to not support recursion lends itself well to smart contracts, since contracts need to terminate in accordance with the parameters of execution. Thus, Turing-incompleteness is a highly sought-after property in smart contract development. While Turing-completeness is a property of the programming language used for smart contracts, Byzantine-fault tolerance is a property of the underlying consensus algorithms used by blockchain nodes to process transactions for the blockchain. Fault-tolerance is another key aspect of the consensus algorithms that drive the distributed nature of blockchain. Since nodes in the chain approve the inclusion of a transaction for a block by means of the outcome of the consensus algorithm, the consensus process needs to be fault-tolerant. Fault-tolerance falls under two categories: crash fault-tolerance, where a node fails to respond and Byzantine fault tolerance, where nodes fraudulently tamper with the data to interfere with the consensus process. Another issue surrounding decentralized contracts is their alignment with the legal frameworks of contract law, that often vary between various jurisdictions and have varying interpretations based on the type of programming language used to code the smart contract, the application of the smart contract, and even its claim of being a contract as defined by contract theory. Thus, the decentralized nature of blockchain adds increasing complexity to dispute resolution and legality surrounding smart contracts.

6. Conclusions

DeFi holds the potential for revolutionizing existing solutions in banking and finance. By rethinking the traditional models of markets, lending, and centralized governance, DeFi offers the promise of widespread access to capital and novel mechanisms for financial transactions. We highlighted flash loans as one such novel mechanism, that distinguishes itself from traditional loans with the triad of zero collateral, reversal of the funds at the end of the transaction, and the ability to potentially borrow the entire liquidity pool and make multiple transactions with the borrowed funds for arbitrage opportunities. However, flash loans have also been subject to attacks of various kinds, ranging from flaws in the smart contracts, protocols, and the oracles. As DeFi applications and adoption continue to rise and they get increasingly integrated with the global financial market, it is important to consider the role of governance models, regulation, and security in the design and development of DeFi offerings.

Table I: A comparison of DEX swap exchanges

Parameters	PancakeSwap	Uniswap	BurgerSwap	BakerySwap
Transaction Cost	0.2%	.3%-1%	0.3%	0.3%
Runs on	BSC	Ethereum	BSC	BSC
Price	$14.75	$19.61	$4.03	$1.96
Market Cap	$2.9 Bil	$11 Bil	$51 Mil	$353 Mil
Number of Circulating Native Tokens	201 M CAKE	520 M UNI	21 M BURGER	180 M BAKE

References

Al-Breiki, H., Rehman, M. H. U., Salah, K., & Svetinovic, D. (2020). Trustworthy blockchain oracles: Review, comparison, and open research challenges. *IEEE Access*, *8*, 85675-85685.

Aigner, A. A., & Dhaliwal, G. (2021). UNISWAP: Impermanent Loss and Risk Profile of a Liquidity Provider. *arXiv preprint arXiv:2106.14404*.

Angeris, G., Kao, H. T., Chiang, R., Noyes, C., & Chitra, T. (2019). An analysis of Uniswap markets. *arXiv preprint arXiv:1911.03380*.

Back, A. (2002). Hashcash - a denial of service counter-measure. Available at : http://www.hashcash.org/papers/hashcash.pdf

N., Frame, W. S., & Ioannidou, V. (2016). Reexamining the empirical relation between loan risk and collateral: The roles of collateral liquidity and types. Journal of Financial Intermediation, 26, 28-46.

Binance Academy. (2020, October 5). *An Introduction to Binance Smart Chain (BSC)*. Binance Academy. https://academy.binance.com/en/articles/an-introduction-to-binance-smart-chain-bsc

Binance Smart Chain: A Parallel Binance Chain to Enable Smart Contracts. Binance Chain. (2020, April 17). https://dexin.bnbstatic.com/static/Whitepaper_%20Binance%20Smart%20Chain.pdf.

Blazy, R., & Weill, L. (2013). Why do banks ask for collateral in SME lending? Applied Financial Economics, 23(13), 1109-1122.

Bitcoin News and Reports (2021). DeFi Exploits in 2021: A Compilation of DeFi Crimes. Available at: https://bitcoinnewsandreports.com/defi-exploits-in-2021-a-compilation-of-defi-crimes/

Bourgi, S. (2021). Spartan Protocol exploit results in loss of $30M. Cointelegraph. Available at: https://cointelegraph.com/news/spartan-protocol-exploit-results-in-loss-of-30m

Browne, R. (2021). Suspected hacker behind $600 million Poly Network crypto heist did it ‘for fun’. CNBC Tech. Available at: https://www.cnbc.com/2021/08/12/poly-network-hacker-behind-600-million-crypto-heist-did-it-for-fun.html

What is MakerDAO?: Research & Fundamentals. Bitcoin Suisse. (2020, December 20). https://www.bitcoinsuisse.com/fundamentals/what-is-makerdao.

Boot, A. W., & Thakor, A. V. (1994). Moral hazard and secured lending in an infinitely repeated credit market game. International economic review, 899-920.

Caldarelli, G., & Ellul, J. (2021). The Blockchain Oracle Problem in Decentralized Finance—A Multivocal Approach. Applied Sciences, 11(16), 7572.

Chan, Y. S., & Kanatas, G. (1985). Asymmetric valuations and the role of collateral in loan agreements. Journal of money, credit and banking, 17(1), 84-95.

Chen, Y., & Bellavitis, C. (2019). Decentralized finance: Blockchain technology and the quest for an open financial system. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3418557

Chen, Y. R., Leu, J. S., Huang, S. A., Wang, J. T., & Takada, J. I. (2021). Predicting Default Risk on Peer-to-Peer Lending Imbalanced Datasets. IEEE Access, 9, 73103-73109.

Cheung, S., & Sundaresan, S. (2006). Lending Without Access to Collateral A Theory of Micro-Loan Borrowing Rates. Work.

Chohan, U. W. (2021). Decentralized Finance (DeFi): An Emergent Alternative Financial Architecture. Critical Blockchain Research Initiative (CBRI) Working Papers.

Collins, R. (1980). Weber's last theory of capitalism: a systematization. American Sociological Review, 925-942.

Conley, J. P. (2017). Blockchain and the economics of crypto-tokens and initial coin offerings. *Vanderbilt University Department of economics working papers*, (17-00008).

Crawley, J. (2021). Flash Loan Attack Causes DeFi Token Bunny to Crash Over 95%. CoinDesk Markets. Available at: https://www.coindesk.com/markets/2021/05/20/flash-loan-attack-causes-defi-token-bunny-to-crash-over-95/

Elahi, Q. K., & Rahman, L. M. (2006). Micro-credit and micro-finance: functional and conceptual differences. Development in Practice, 16(5), 476-483.

Fenu, G., Marchesi, L., Marchesi, M., & Tonelli, R. (2018, March). The ICO phenomenon and its relationships with ethereum smart contract environment. In *2018 International Workshop on Blockchain Oriented Software Engineering (IWBOSE)* (pp. 26-32). IEEE.

Hanson, R. (2003). Combinatorial information market design. Information Systems Frontiers, 5(1):107–119.

Heasman, W. (2020). Are the BZx Flash Loan Attacks Signaling the End of DeFi?. Cointelegraph. Available at: https://cointelegraph.com/news/are-the-bzx-flash-loan-attacks-signaling-the-end-of-defi

Jarboe, K. P., & Ellis, I. (2010). Intangible assets: Innovative financing for innovation. Issues in Science and Technology, 26(2), 75-80.

Jimenez, G., Salas, V., & Saurina, J. (2006). Determinants of collateral. Journal of Financial Economics, 81(2), 255-281.

John P. Caskey. 1997. Lower Income Americans, Higher Cost Financial Services. Filene Research Institute and the Center for Credit Union Research, Madison, WI. (Monograph)

Kahraman, E. (2021). Binance Smart Chain-based DeFi platform suffers $3M flash loan attack. Cointelegraph. Available at: https://cointelegraph.com/news/binance-smart-chain-based-defi-platform-suffers-3m-flash-loan-attack

Khatri, Y. (2019). There are now at least 11 blockchain unicorns with over $1B valuation, including Binance, Ripple and Coinbase. The Block.
https://www.theblockcrypto.com/post/44277/there-are-now-11-blockchain-unicorns-in-the-world-with-over-1b-valuation-including-binance-ripple-and-coinbase

Laboure, M. and Braunstein, J. (2021). The Social Consequences of Decentralized Finance. The Journal of Global Policy. Wiley.

Loumioti, M. (2012). The use of intangible assets as loan collateral. Available at SSRN 1748675.

Neuder, M., Rao, R., Moroz, D. J., & Parkes, D. C. (2021). Strategic Liquidity Provision in Uniswap v3. *arXiv preprint arXiv:2106.12033*.

Niinimäki, J. P. (2011). Nominal and true cost of loan collateral. Journal of Banking & Finance, 35(10), 2782-2790.

Othman, A., Pennock, D. M., Reeves, D. M., & Sandholm, T. (2013). A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation (TEAC)*, *1*(3), 1-25.

Popescu, A. D. (2020). Decentralized finance (defi)–the lego of finance. Social Sciences and Education Research Review, 7(1), 321-349.

Qin, K., Zhou, L., Livshits, B., & Gervais, A. (2021). Attacking the DeFi ecosystem with flash loans for fun and profit. In International Conference on Financial Cryptography and Data Security (pp. 3-32). Springer, Berlin, Heidelberg.

Rajan, R. G. (1992). Insiders and outsiders: The choice between informed and arm's-length debt. *The Journal of Finance*, *47*(4), 1367-1400.

Redman, J. (2021). BSC Defi Protocol Burgerswap Loses $7.2 Million from a Flash Loan Attack. Bitcoin News. Available at: https://news.bitcoin.com/bsc-defi-protocol-burgerswap-loses-7-2-million-from-a-flash-loan-attack/

Rotfleisch, D. (2021). Tax Implications Of Cryptocurrency Flash Loans. Available at: https://www.mondaq.com/canada/fin-tech/1085052/tax-implications-of-cryptocurrency-flash-loans

Salami, I. (2020). Decentralised Finance: The Case for a Holistic Approach to Regulating the Crypto Industry. Journal of International Banking and Financial Law, 35(7), 496-499.

Samieifar, S., & Baur, D. G. (2021). Read me if you can! An analysis of ICO white papers. *Finance Research Letters*, *38*, 101427.

Thurman, A. (2021). Rari Capital falls victim to $11 million exploit. Cointelegraph. Available at: https://cointelegraph.com/news/rari-capital-falls-victim-to-11-million-exploit

Vermaak, W. (2021). What Are Flash Loan Attacks?. Coinmarketcap. Available at: https://coinmarketcap.com/alexandria/article/what-are-flash-loan-attacks

Wang, D., Wu, S., Lin, Z., Wu, L., Yuan, X., Zhou, Y., ... & Ren, K. (2021, May). Towards A First Step to Understand Flash Loan and Its Applications in DeFi Ecosystem. In *Proceedings of the Ninth International Workshop on Security in Blockchain and Cloud Computing* (pp. 23-28).

Werhane, P. H. (2002). Moral imagination and systems thinking. Journal of Business Ethics, 38(1), 33-42.

Wright, A. (2020). The Growth and Regulatory Challenges of Decentralized Finance. Report of the TAC Virtual Currency Subcommittee. Available at: https://www.cftc.gov/media/5471/TAC121420_GrowthRegulatoryChallengesDecentralizedFinance/download

Young, M. (2021). Impossible Finance Loses $500,000 in Latest DeFi Flash Loan Attack. Yahoo. Available at: https://www.yahoo.com/now/impossible-finance-loses-500-000-064341634.html

Young, M. and Baird, K. (2021). DeFi Protocol bEarn Suffers $11M Flash Loan Attack. Bein Crypto Technology News Report. Available at: https://beincrypto.com/defi-protocol-bearn-suffers-11m-flash-loan-attack/

Zetzsche, D., Arner, D., Buckley, R. (2020. Decentralized Finance, Journal of Financial Regulation, Volume 6, Issue 2, 20 September 2020, pages 172–203, https://doi.org/10.1093/jfr/fjaa010