

Aaro Capital

An Introduction to Distributed Ledger Technology

Disclaimer

The material provided in this document is being provided for general informational purposes. Aaro Capital Limited does not provide, and does not hold itself out as providing, investment advice and the information provided in this document should not be relied upon or form the basis of any investment decision nor for the potential suitability of any particular investment. The figures shown in this presentation refer to the past or are provided as examples only. Past performance is not reliable indicator of future results.

This document may contain information about cryptoassets. Cryptoassets are at a developmental stage and anyone thinking about investing into these types of assets should be cautious and take appropriate advice in relation to the risks associated with these assets including (without limitation) volatility, total capital loss, and lack of regulation over certain market participants. While the directors of Aaro Capital Limited have used their reasonable endeavours to ensure the accuracy of the information contained in this document, neither Aaro Capital Limited nor its directors give any warranty or guarantee as to the accuracy and completeness of such information.

Please be sure to consult your own appropriately qualified financial advisor when making decisions regarding your own investments.

Acknowledgements

Dr. Spyros Galanis and Peter Habermacher would like to thank Ankush Jain, Dr. Giovanni Ricco, Dr. Daniele Bianchi, Dr. Jerome Rousselot and Oscar Pacey for their input on this paper.

Executive Summary

In this paper, we review the economics behind distributed ledger technology (DLT) and related concepts, including Bitcoin, blockchain and cryptocurrencies.

Markets that operate efficiently have systems of maintaining records on transactions that have taken place. Importantly, buyers and sellers are required to trust that this information is kept safe and updated correctly, so that they have a path of recourse should a dispute arise. For example, a buyer on eBay has to trust that the network will register their transfer of money and ultimately send the goods they purchased, despite never physically meeting the seller.

Until recently, updating and maintaining records of transactions could only be performed by trusted intermediaries, such as banks, firms or governments, operating in an environment with strong institutions. The unintended consequence is that these intermediaries may obtain market power, which they may abuse. Market power can arise when a firm undertakes a large investment to become a trusted intermediary, via an extended period of building reputation and, in some cases, becoming regulated.

DLT can alleviate this issue because the users of the distributed ledger can also be the owners. More importantly, breaking the monopoly over the ownership of the ledger has the potential to create more efficient and trusted systems of disseminating information. DLT promotes trust among market participants because all elements of a transaction that are recorded on a ledger can be reliably and directly verified at low cost, by any participant.

DLT also has the potential to solve issues of trust that arise from the “hold-up” problem. For example, when a contributor (firm A) deposits its data on a database controlled (fully or partially) by an administrator (firm B), it makes a substantial investment that has little value outside of the relationship between these two firms. However, contracts are almost always incomplete, meaning that unforeseen contingencies might arise in the future, so that the two parties need to renegotiate their relationship. At this point, firm A is held up by firm B and may therefore be forced to accept worse terms during the renegotiation. DLT alleviates the hold-up problem, mainly because ownership of the ledger is shared, so there is no single owner who could abuse their market power at a future date.

Another manifestation of the interplay between trust and market power is the familiar “chicken and egg” problem that any new network faces. A network’s value increases as more participants (users, developers, investors) join, but their participation depends on the network already being valuable. Traditional revenue models solve this problem of trust by granting early participants (usually investors) property rights over the network, so that if it becomes valuable, they get rewarded. However, in many cases the unintended consequence is that these participants also gain excessive market power. DLT has the potential of solving this issue of trust, while limiting the market power gained by platform contributors and early adopters. This is achieved by issuing a token on the network, which is earned by participants (users, developers and investors) through various forms of contributions to the network. The token may generate economic value for its holders through mechanisms such as network voting rights, or as a means of payment between network participants. With the correct token design, the incentives of network participants can be aligned.

A distributed ledger is just one of many possible database structures. The most important properties of a database are the control system and the execution architecture. There are sliding scales for each of these properties from highly centralised to decentralised, each with their own advantages and disadvantages. For both, centralisation has large benefits in terms of efficiency, and therefore centrally controlled and stored databases are the most common. However, centralised databases are far less robust than their distributed counterparts.

A distributed ledger is a specific type of a distributed database, based on and verified by the mathematical properties of cryptography. Introducing a cryptography-based data structure makes the ledger immutable and append only. However, immutable, read-only databases are nothing new and can easily be created by changing the write permissions of a database. The key innovation is cryptography-based data structure, which for the first time enables digital scarcity.

Cryptography makes distributed ledgers far more suited for instances where trust between participants is an issue, and a database is to be governed in a decentralised manner. There is no need to trust other participants on the ledger, and this is its key advantage over a traditional distributed database.

There are different types of distributed ledgers, the most common being the blockchain. A blockchain consists of possibly several chains of blocks. Each block contains pieces of information, such as financial transactions. The order of blocks matters. Although several chains may coexist temporarily, there is consensus on the one that everyone follows and updates according to some rule defined by the blockchain protocol.

While distributed ledgers are peer-to-peer in terms of how the data is stored, the control over the ledger may be centralised or decentralised.

The term permissioned ledgers typically refers to a jointly controlled and maintained ledger, with a controlled user base and a small number of semi-trusted ledger writers/controllers. This allows for greater ledger control, greater customisation and does not require a cryptocurrency or sybil resistance mechanism to align incentives. Permissioned ledgers also use a different consensus algorithm than most permissionless ledgers. These ledgers are favoured by enterprises for business-to-business transactions. Although permissionless DLTs can achieve transparency and decentralisation, enterprises often value privacy and control, together with the fast processing and finality of transactions.

DLT is advantageous over traditional centralised or shared databases in low trust environments. Examples include:

1. International Remittances, Cross Border Transfers and Clearance of Payments
2. Trade Finance
3. Supply Chains
4. Insurance
5. Healthcare

In contrast, public, or permissionless ledgers, are open to everyone. Anyone can run a network node to verify their own copy of the ledger; they may choose to extend the ledger by competing in mining for blocks, as well as develop the open source code on which it runs. Permissionless ledgers use game theory to align the incentives of all users. The most famous example of a permissionless distributed ledger is Bitcoin. Bitcoin (BTC) is the first, most well-known and largest cryptocurrency, implementing a permissionless and distributed blockchain. The cryptocurrency's primary function on top of the bitcoin ledger is to act as an incentive and coordination mechanism that prevents attacks that corrupt the data stored in the ledger.

We review how a transaction happens on the Bitcoin ledger. Suppose that Ann buys a pizza from Bob for 1 Bitcoin (BTC). The transaction records the transfer of 1 BTC between the two public addresses, one for Bob and one for Ann. Ann's public address acts as her account and is visible to everyone. However, she herself is not visible, and is the only person who has access to her account's private key, or password. Ann signs the transaction, by proving that she controls the public address from which the 1 BTC is transferred to Bob. Using properties of cryptography, she can prove that she controls the public address using her private key, without revealing her private key. The transaction, together with her signature, is broadcast to the network of miners.

On average every 10 minutes, a collection of new transactions is validated. They are written into a block by a miner who adds a reward of 12.5 BTC to themselves, together with transaction fees. These 12.5 BTC are newly created and are recorded on the ledger for the first time. The miner is chosen at random, using the Proof-of-Work protocol, which specifies that the winner who has solved a cryptographic puzzle first gets to propose the next block (solving this problem relies on luck and computational power). The lucky miner broadcasts their solution, together with the new block. All other miners verify the solution and append the block to the blockchain. Although finding the solution is difficult and costly in terms of computational power and electricity, verifying that the solution is correct is instant. The reward halves periodically, so the rate of creating new BTC converges to zero over time. Approximately 21 million BTC are scheduled to be created by 2140.

There are also other permissionless distributed ledgers, each with their own distinctive properties. Ethereum offers a virtual machine, where smart contracts and decentralised applications (dApps) can be implemented. It is intended to become the world's distributed computer, where programmers can concentrate on building dApps for a variety of uses on top of an existing distributed ledger infrastructure.

Ripple achieves scalability and speed by avoiding the use of the Proof-of-Work protocol. Instead, it uses a low-latency Byzantine agreement protocol, which can reach consensus without full agreement of all nodes. Moreover, its cryptocurrency, XRP, was created at inception, instead of being created with every block. The intended use of XRP is as a bridge currency that facilitates foreign exchange and business-to-business payments.

Zcash is a cryptocurrency focused on the privacy of transactions. To achieve this, Zcash uses zero-knowledge proofs and two types of addresses: private (z-addresses) and transparent (t-addresses), where the latter is similar to the public addresses of Bitcoin. A transaction can be Z-to-Z, meaning that it is recorded on the public blockchain and known to have occurred, however the amount, the fees and the addresses are encrypted and private. A T-to-T transaction is similar to a transaction recorded in Bitcoin, where the addresses, the fees and the amount are public.

Delving deeper into exactly how permissionless ledgers work, one of the most important issues in the design of a blockchain is how consensus on the correct state of the ledger is achieved, as well as who is going to write the next block. The easiest way of choosing the writer of the next block is to randomly pick one participant. However, this leads to the possibility of a "Sybil attack", where a participant creates multiple selves (e.g. multiple IP addresses) in order to increase their probability of selection and the payoff that they will receive. If a participant greatly increases their probability of selection, they can also control the ledger. The Proof-of-Work and Proof-of-Stake protocols are solutions to this problem and are therefore called "Sybil resistance mechanisms". They generate scarcity of resources, making it increasingly difficult and expensive for a participant to create multiple selves. The Proof-of-Work protocol achieves resistance by selecting the participant (miner) who can first solve a difficult (and costly in terms of computation) problem. Proof-of-Stake specifies that the probability of selection is proportional to the miner's stake of coins, which are by construction scarce and cannot be replicated.

In both instances, the cryptocurrency acts as compensation for an upfront cost (electricity and mining equipment or buying coins to stake). The more coins or mining equipment a participant has, the more they have to lose if the cryptocurrency were to decline in value. Thus, as the influence of a participant increases, the more incentivised they are to maximise the value of the cryptocurrency via the network's benefit to users. Therefore, a cryptocurrency on top of a distributed ledger is primarily an incentive and coordination mechanism in the absence of any trusted controlling entity.

The other main issue is reaching consensus on which branch of the blockchain the new block of information is to be attached. The two branches might have been created because of lack of communication and latency, or because some malicious participants alter the information in previous blocks and want to make their branch the correct one. The most common consensus algorithm for permissionless ledgers is called Nakamoto Consensus.

There are still many limitations of distributed ledgers preventing it from achieving mainstream adoption - it is important to remember that DLT is at the early stages of research and development. In order to understand the evolution of this technology, we draw analogies from previous technology market cycles, such as in hardware (1950-1970), software (1970-1990) and networks (1990-2010). Expansion is first typically driven by open standards and decreasing costs. Then there is a phase of consolidation, where winners build proprietary systems on top of these open standards, stifling competition. Finally, there is decentralisation through the development of open source alternatives, in order to escape the platforms of incumbent firms and their high fees. DLT is still at the early stages of expansion, where open standards are being developed in order to decrease costs.

Note that there are still obstacles in terms of how well the technology can scale and expand. The scalability trilemma specifies that it is very difficult (if not impossible) to achieve the following three desirable objectives simultaneously: safety, scalability and decentralization. Safety refers to whether a blockchain can withstand a malicious attack, one that aims to corrupt or reverse recorded transactions. Scalability is measured by the number of transactions per unit of time that the system can perform. Decentralisation of block production (DBP) is defined as the number of independent block producers and how easy it is for a new participant to become a block producer. There are many different ledger designs which achieve two out of three objectives satisfactorily, but not all three. For example, Bitcoin achieves safety and decentralisation, but not scalability.

Layer 2 solutions provide an alternative way to solve for the scalability trilemma, particularly in terms of achieving greater scalability on various dimensions. These protocol projects work by performing some computations off-chain, while still anchoring to the main blockchain to maintain security and trustlessness.

An important example of a Layer 2 solution is the concept of a sidechain. A sidechain is a separate blockchain that attaches to the main blockchain. The two chains communicate (sometimes in predetermined intervals), so that tokens from the mainchain are transferred to the sidechain. When the transfer is complete, computations can be performed on the sidechain. When the computations are complete, the tokens are transferred back to the main blockchain. The mainchain only records the initial and the final states, whereas the sidechain records all intermediate states (e.g. intermediate transactions between two parties). If a dispute on the sidechain arises that cannot be resolved there, it is resolved in the mainchain by reinstating the initial state and punishing participants or redoing calculations on the mainchain (which is costly). This acts as an incentive for participants to be truthful and cooperative.

Another example of a Layer 2 solution is the Lightning Network. It is a payment network on top of the Bitcoin blockchain, enabling two users to establish a bidirectional private payment channel and then perform many transactions between them. Transactions can settle much faster at a lower cost, since users only need to record their initial and final transactions on the blockchain. This is done by using a smart contract, which is essentially a balance sheet. When all transactions are complete, the connection terminates and the amounts on the balance sheet are recorded in the blockchain. When the network expands, users are not required to establish a direct channel with each person they want to transact with, as the Lightning Network can find an indirect path in order to establish a connection.

One of the main criticisms against Bitcoin and cryptocurrencies in general is that they are designed to facilitate illegal behaviour: they allow for pseudonymous transactions which do not reveal the identity of transacting parties. This is no longer true. A report by Chainalysis shows that the share of value in BTC sent to darknet markets has declined from 7% in 2012 to less than 1% in 2018. There are two reasons for this. First, the regulation has been updated, and law enforcement authorities have started to act against these cases. Second, the design of the blockchain, where transactions are public but pseudonymous, helps rather than hinders authorities in their effort to prosecute illicit uses.

Contents

Disclaimer	i
Acknowledgements	ii
Executive Summary	iii
Contents	vii
1 Transactions, Trust and Market Power	1
2 Distributed Ledgers	7
2.1 Database Structures	7
2.2 Databases vs Distributed Ledgers	8
2.3 Distributed Ledger Structures	9
2.4 Blockchains	10
2.5 Public, Private and Permissioned Ledgers	11
3 Permissioned Distributed Ledgers	14
3.1 Distributed Ledger Platforms for Enterprises	14
3.2 International Remittances, Cross Border Transfers and Clearance of Payments	14
3.3 Trade Finance	15
3.4 Supply Chains	16
3.5 Insurance	17
3.6 Healthcare	18
4 Permissionless Distributed Ledgers	20
4.1 Bitcoin	20
4.2 Ethereum	22
4.3 Ripple	24
4.4 Zcash	25
5 Consensus	27
5.1 Proof-of-Work	27
5.2 Proof-of-Stake	28
5.3 Nakamoto Consensus	29
5.4 Security of Permissionless Ledgers and Traditional Databases	31
6 The DLT Market Cycle	33
7 Design Issues and Solutions	36
7.1 The Scalability Trilemma	36
7.2 Layer 2 Solutions	37
7.2.1 Sidechains	37
7.2.2 Lightning Network	38
7.3 Illicit Uses	39

1 Transactions, Trust and Market Power

Markets that operate efficiently have systems of maintaining records on transactions that have taken place. buyers and sellers are required to trust that this information is kept safe and updated correctly, so that they have a path of recourse should a dispute arise. For example, a buyer on eBay has to trust that the network will register their transfer of money and ultimately send the goods they purchased, despite never physically meeting the seller.

Until recently, updating and maintaining records of transactions could only be performed by trusted intermediaries, such as banks, firms or governments, operating in an environment with strong institutions. The unintended consequence is that these intermediaries may obtain market power, which they may abuse. Market power can arise when a firm undertakes a large investment to become a trusted intermediary, via an extended period of building reputation and, in some cases, becoming regulated.¹ Once trust is established within a network, it often becomes difficult for participants to create or join a different network, thus reinforcing the dominant position of the trusted intermediary. Market power can also arise due the intermediary's ability to use the information it possesses about past transactions to gain valuable insights about the network's users and their preferences. However, the incentives of these trusted intermediaries are not necessarily aligned with those of the market participants that provide the information, thus creating a principal-agent problem. Although the users (principals) provide the information that is crucial to sustain the market, the intermediaries (agents) could potentially use this information in ways that are against their best interests (e.g. by raising fees).

Over the years, there have been many cases of information abuse by trusted intermediaries. Examples include the €2.4bn fine by the European Commission on Google, who abused their market dominance to the benefit of its own comparison-shopping service.² Facebook has been fined the maximum amount possible by a UK regulator for failing to protect its users' personal information and not being transparent on how their data is used.³ Global banks have been fined \$321bn since the financial crisis.⁴ Moreover, data breaches at various multinationals indicate inadequate procedures in place to safeguard the personal data of customers.⁵

Distributed ledger technology (DLT), which enables the decentralised creation and maintenance of ledgers, has the potential to revolutionize the way we store and distribute information - and therefore could change the way markets operate, by realigning the incentives of all market participants. Unlike centralised or shared databases, the users of distributed ledgers can also be the owners. Breaking the monopoly over the ownership of a ledger has the potential not only of redistributing wealth, but also creating more efficient and trusted systems of disseminating information, thus widening participation and intensifying competition.

¹ Reputation building via marketing activities often falls under rent-seeking activities which can hamper the economic efficiency of an economy.

² For more information, see: http://europa.eu/rapid/press-release_IP-17-1784_en.htm.

³ For more information, see: <https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2018/10/facebook-issued-with-maximum-500-000-fine/>.

⁴ For more information, see <https://www.bloomberg.com/news/articles/2017-03-02/world-s-biggest-banks-fined-321-billion-since-financial-crisis>.

⁵ For more information, see: https://en.wikipedia.org/wiki/List_of_data_breaches

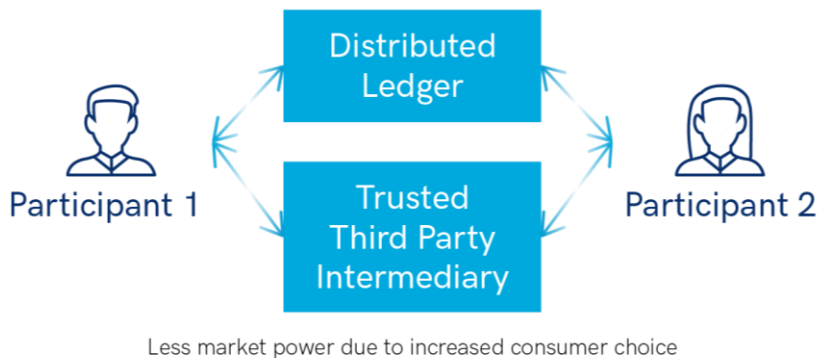
Figure 1: Moderating market power via decentralised ledgers

Without Distributed Ledgers:



With Distributed Ledgers:

An alternative to trusted third party intermediary if fees and terms are unfavourable relative to trustless peer-to-peer transactions



Source: Aaro Capital Research

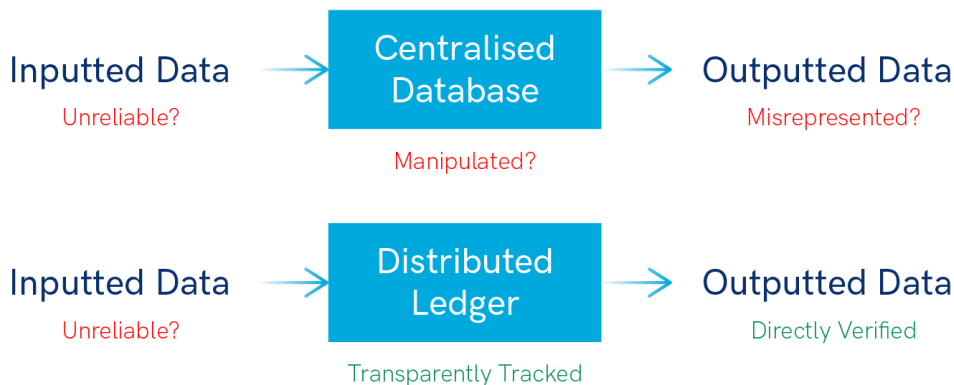
DLT promotes trust among market participants because all elements of a transaction that are recorded on a ledger (on-chain) can be reliably verified at low cost.⁶ Hence, there is no longer the need to rely on trusted intermediaries - potentially reducing the market power they may hold.⁷ The low cost of verification ensures that information entered on the ledger can be transparently tracked, thus guaranteeing its integrity. Moreover, any user can read the ledger directly, without relying on third parties who may misrepresent the information. However, there is still the issue of verifying elements of a transaction that exist outside the ledger. For example, a clause in a transaction may be triggered only when it is verified that a “real world” event has occurred. In those cases, a trusted intermediary may still be needed.⁸

⁶ A more detailed discussion on the cost savings of verification via DLT can be found at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2874598.

⁷ In competition theory, simply the threat of a credible outside option is enough to dissipate all market power. However, this may not happen in the real world due to various reasons.

⁸ Distributed ledgers do not address this issue directly, but there are projects which aim to mitigate this problem, by providing solutions that vary depending on the application.

Figure 2: Trust issues with centralised databases

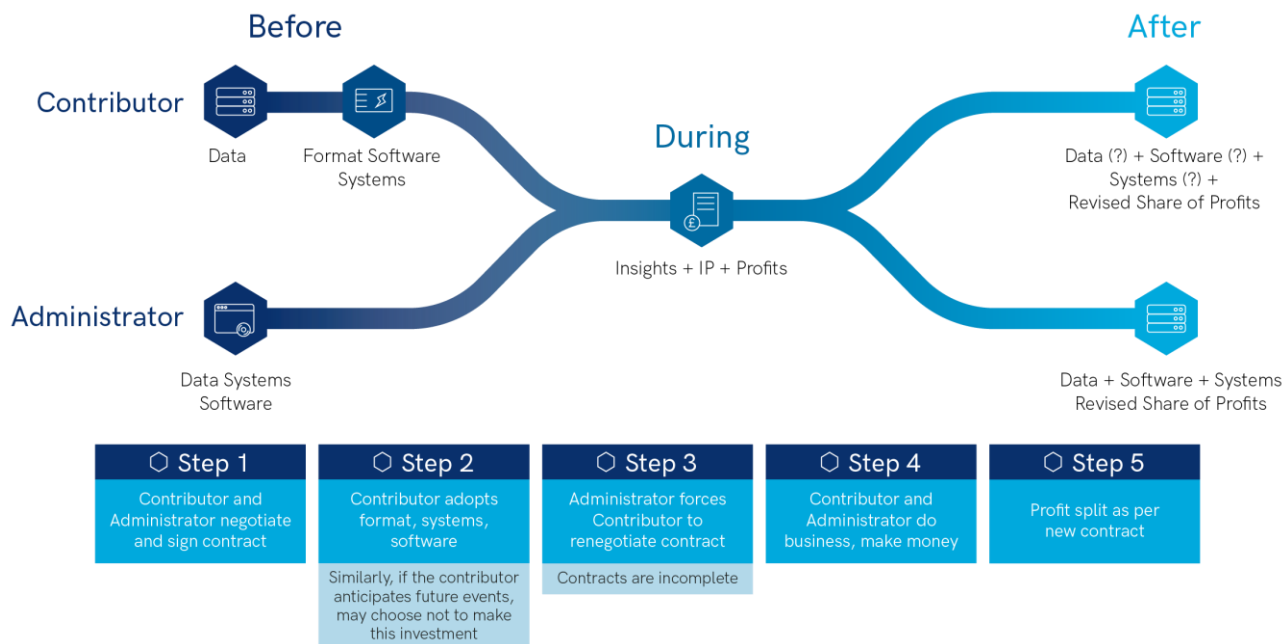


Source: Aaro Capital Research

DLT also has the potential to solve issues of trust that arise from the “hold-up” problem, as outlined in Figure 3.⁹ When a contributor (firm A) deposits its data on a database controlled (fully or partially) by an administrator (firm B), it makes a substantial investment that has little value outside of the relationship between these two firms. Firm A may have to adopt new hardware and/or software to ensure that their data is mutually compatible. Furthermore, staff at firm A will need training on how to use the new systems. If the two firms were able to negotiate a complete contract on how this data is used, that foresees all possible future contingencies, then this investment would not be a problem, because at the beginning of the relationship both parties have equal bargaining power. However, contracts are almost always incomplete, meaning that unforeseen contingencies might arise in the future, such that the two parties need to renegotiate their relationship. At that point, firm B has a strong bargaining position, because if firm A walks away, it loses its data and its initial investment. In other words, firm A is held up by firm B, and may therefore be forced to accept worse terms during the renegotiation. Moreover, contract negotiations and renegotiations are time-consuming and expensive.

⁹ A more detailed discussion is provided at <https://docsend.com/view/zbq3bud>.

Figure 3: Hold up problem for Consortia



Source: Prysm Group

The hold-up problem is compounded in the case of databases, due to the complexity of property rights on data. First, it is common for the system administrator, firm B, to be the ultimate owner of the data, so that firm A loses control of the data it uploads on the database. Second, this data may be stored on a third-party who is using firm A as a trusted intermediary. Third, the ease of copying data weakens the bargaining power of firm A further. While firm A can threaten to leave the shared database and delete their data, it is hard to restrict firm B's access to it, because it may have its own copy. Moreover, it is difficult for firm B to unlearn the knowledge it has gained by analysing A's data. Because data can be used in many different ways, it is almost impossible to write a complete contract that specifies all future contingencies.¹⁰ Finally, if firm A chooses to leave the shared database, it may be left with data that cannot be read without access to the software provided by firm B.

Due to the reasons above, a market participant may anticipate the hold-up problem and choose not to upload data to a shared database, thus reducing cooperation and economic value creation. For example, hospitals in the USA have created a system of fragmented digital data silos, due to the technical and, in particular, economic issues that shared databases create. This is in contradiction to the US HITECH Act of 2009, which envisaged seamless electronics transmission of medical data.¹¹

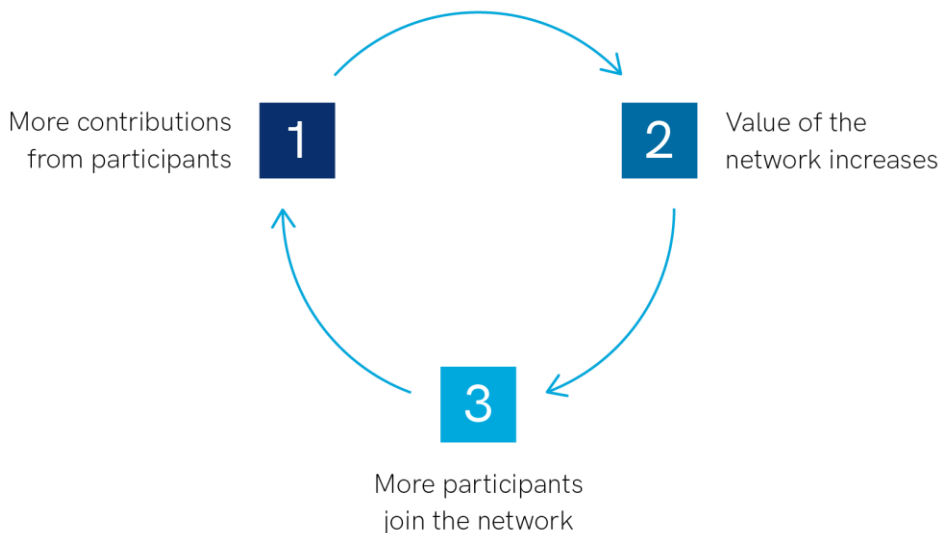
DLT alleviates the hold-up problem, mainly because ownership of the ledger is shared, so there is no single owner who could abuse their market power at a future date. Moreover, smart contracts allow participants to write programs that are executed automatically when certain events occur, thus making enforcement easier. More importantly, each contributor has greater control over the data that they share in the database, as they can grant and revoke rights on who is able to read this information, at any time.

¹⁰ The EU introduced GDPR to try and solve the issues faced when sharing and processing data.

¹¹ For more information on the US HITECH Act of 2009, see: <https://www.hhs.gov/hipaa/for-professionals/special-topics/hitech-act-enforcement-interim-final-rule/index.html>.

Another manifestation of the interplay between trust and market power is the familiar “chicken and egg” problem that any new network faces. A network’s value increases as more participants (users, developers, investors) join, but their participation depends on the network already being valuable. Moreover, early participants face a free-riding problem because they contribute their resources to the success of the network. If a network succeeds, then all participants benefit, even those who did not contribute. If early participants cannot trust that the network will proportionally reward them when it succeeds, they will have fewer incentives to contribute, and the network will not be as valuable.

Figure 4: The “chicken and egg” issue faced by new networks



Source: Aaro Capital Research

Traditional revenue models solve this problem of trust by granting early participants (usually investors) property rights over the network, so that if it becomes valuable, they get rewarded. They are then incentivised to contribute the resources required for the network to succeed. However, in many cases the unintended consequence is that these participants also gain excessive market power, which they often use to the detriment of other participants of the network.

DLT has the potential of solving this issue of trust, while limiting the market power gained by platform contributors and early adopters.¹² This is achieved by issuing a token on the network, which is earned by participants (users, developers and investors) through various forms of contributions to the network. Tokens generate economic value to holders through mechanisms such as network voting rights, or as a means of payment between network participants.¹³ If the network succeeds, the value of the token increases and participants get rewarded, depending on their individual contribution. As there are many contributors, it is much harder for an individual participant to gain meaningful market power that may be abused later. Further, the larger the holdings of the token and thus the greater the influence a participant has over the network, the greater incentive they have to maximise the value of the token - typically achieved via maximising the value of the network to its users. Thus, the incentives

¹² A detailed discussion on how DLT can solve the “tragedy of the commons” problem faced by networks be found at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2874598.

¹³ Token economic design is critical to the value of the token. There must be economic benefit for holding the token beyond speculation (where many tokens failed in 2017/18). This is covered in more detail in the paper “*An Introduction to Web 3.0*”.

of network participants should be aligned.¹⁴ Further, if a large token holder does act against the best interests of the network, it is easy for users create a new parallel network via a hard fork, which removes the problem user.¹⁵ In other words, the network effects of platforms, such as Facebook, can be disconnected from the data and protocol layers where the market power lies, thus potentially solving the competition issues that data / internet giants currently pose.¹⁶

Table 1: The different layers of a platform

	Traditional	Permissionless Ledger
Data Layer	Platform owners typically own and control user data	Users own and control their data
Network Layer	Network effects usually lead to market consolidation	Network effects usually lead to market consolidation
Protocol Layer	Platform owners typically own and control platform protocols	Users typically own and control platform protocols

Source: Aaro Capital Research

To summarise, DLT has the potential to revolutionise the way that markets operate, by increasing trust between market participants and enabling them to create more valuable networks. At the same time, DLT can mitigate several of the market failures which are associated with the increased market power of trusted intermediaries. These include the inefficiencies generated by incomplete contracts and the hold-up problem, free riding when contributing to a network, and misaligned incentives between those who control the network and those who contribute to it.

¹⁴ There are scenarios where there may be misaligned incentives. For example, major holders of a larger network may become holders of a smaller network's tokens in order to destroy it, and therefore force users to move across.

¹⁵ This may have the disadvantage in that it results in smaller competing networks which gain less utility than a combined, larger network. However, it is not in the interest of users to do this arbitrarily, as it will result in the loss of faith of the network and hence a loss in the value of the tokens held.

¹⁶ For a more detailed discussion on the different layers of a distributed ledger, see: <https://www.jbs.cam.ac.uk/faculty-research/centres/alternative-finance/publications/distributed-ledger-technology-systems/#.XWWYH3dFyUk>

2 Distributed Ledgers

2.1 Database Structures

A database is a structured set of digital information, with a unique identification number for each row, and defined rules for the data stored in each column.¹⁷ Historically, the term “ledger” was only used for databases which contained financial transactions. The most important properties of a database are the control system and the execution architecture. There are sliding scales for each of these properties, thus databases come in many forms.

In the dimension of control, the sliding scale varies between completely centralised, where only one entity has read and write permissions, and decentralised, where multiple entities must come to agreement on governance of the database.

Table 2: Centralised vs decentralised administration of databases

	Centralised	Decentralised
Pros	<p>Fewer decision makers in the governance process – quicker and more efficient</p> <p>Most customizable for ease of use</p>	<p>Less reliant on third parties</p> <p>More resilient than a single database administrator</p>
Cons	<p>Only as reliable and resilient as the database administrator</p> <p>Heavy dependence on third party intermediaries</p> <p>Incentives of owner and users may collide</p>	<p>Increasing number of decision makers complicates and slows the governance process, making the system less versatile</p>

Source: Aaro Capital Research

In the dimension of execution architecture, databases fall into three general buckets: centralised databases, decentralised databases and distributed databases, as illustrated Figure 5 below. In centralised databases, a single master copy of the entire database is stored in a single location.¹⁸ Due to its efficiency, scalability and ease of use, it is the dominant type of database. In decentralised databases, data is split between multiple centres. These are commonly used for databases which are too large to store centrally, like those maintained by Google, or when there are multiple data sources feeding into the database. Because some nodes are more important than others and act as “local” centres, bottlenecks may be created. Also, such structures make databases technically challenging to maintain and upgrade. In distributed databases, information is consensually shared among different nodes, dispensing with any centres completely, such that each node in the peer-to-peer network is created equal. The largest challenge of a distributed database is to ensure that these multiple copies are up-to-date and do not conflict with each other. Distributed databases are used in finance, where conflict resolution is a key reason why it still takes days and a high cost for a trade or bank transfer to settle.¹⁹

¹⁷ For more information, see: <https://www.multichain.com/blog/2015/10/private-blockchains-shared-databases/>.

¹⁸ This includes “master and slave” database structures, as one database is used to update other copies.

¹⁹ For more information on multiversion concurrency control (MVCC) mechanisms, see: <https://www.multichain.com/blog/2015/07/bitcoin-vs-blockchain-debate/>.

Figure 5: Database execution structures



Source: Aaro Capital Research

Table 3: A comparison of database execution structures

	Centralised	Decentralised	Distributed
Pros	Highest transaction speed and volume	More resilient to attacks, as there are multiple local centres Data can be stored where it originated, reducing data copying	Most resilient to outside attacks, as there is no single point of failure
Cons	Less robust to attacks as there is a single point of failure May involve copying data from multiple sources into one location Reliant on a trusted third party	Reliant on a trusted third party There are still bottlenecks	Least scalable for transaction speed and volume Least efficient as data is duplicated many times

Source: Aaro Capital Research, Cointelegraph²⁰, Multichain²¹, Ben Morris²²

2.2 Databases vs Distributed Ledgers

A distributed ledger is a specific type of a distributed database, based on and verified by the mathematical properties of cryptography.²³ However, introducing a cryptography-based data structure makes the ledger immutable and append only.^{24,25} To change a data record on a distributed database, one can edit a row of the

²⁰ <https://cointelegraph.com/explained/decentralized-and-distributed-databases-explained>

²¹ <https://www.multichain.com/blog/2015/07/bitcoin-vs-blockchain-debate/>,
<https://www.multichain.com/blog/2016/03/blockchains-vs-centralized-databases/>

²² <https://www.ben-morris.com/a-shared-database-is-still-an-anti-pattern-no-matter-what-the-justification/>

²³ For more information, see <https://hackernoon.com/databases-and-blockchains-the-difference-is-in-their-purpose-and-design-56ba6335778b>.

²⁴ As discussed later, the reversibility of this immutability property depends on various factors.

²⁵ These innovations have also allowed for effective multiversion concurrency control (MVCC), see: <https://www.multichain.com/blog/2015/07/bitcoin-vs-blockchain-debate/>.

database to remove the old data.²⁶ With a distributed ledger, one has to add a new data entry to update an existing piece of data. This feature increases the transparency and traceability of data that the ledger stores. However, immutable, read-only databases are nothing new and can easily be created by changing the write permissions of a database. The key innovation is cryptography-based data structure.

Cryptography makes distributed ledgers far more suited for instances where trust between participants is an issue, and a database is to be governed in a decentralised manner. Distributed ledgers thus inherit properties from both distributed databases and de-centrally governed databases, with several key advantages.

Table 4: An overview of distributed ledgers

	Distributed Ledgers
Pros	Decentralised and secure: most resilient to outside attacks, as there is no single point of failure
	Immutable: information cannot be tampered with or altered
	Transparent: the information is censorship resistant, as it cannot be hidden by third parties retrospectively
	There is no distinction between owners and users
Cons	Least scalable for transaction speed and volume
	Least efficient as data is duplicated many times
	Slowest decision making as users need to reach consensus

Source: Aaro Capital Research, Vince Tabora²⁷

Cryptography allows for the creation of digital scarcity, something that has not previously been attainable.²⁸ Without access to a cryptographic key, it is nearly impossible to make an indistinguishable copy of a cryptoasset.²⁹ For traditional data or databases, it is hard to restrict the copying of data. This is a defining property of distributed ledgers, one that solves market failures such as those caused by incomplete contracts, and allows for the creation of decentralised incentive mechanisms.

2.3 Distributed Ledger Structures

Although distributed ledgers can adopt different structures, as outlined in Figure 6 below, the common feature is that there are multiple independent master copies of the ledger, called network nodes. These nodes share updates to the ledger in a peer-to-peer manner, as previously illustrated in Figure 5. Like traditional databases, the most important properties of a distributed ledger are the execution architecture and the control system.

²⁶ Databases typically have backup functions so there is still some traceability of changes to the database.

²⁷ <https://hackernoon.com/databases-and-blockchains-the-difference-is-in-their-purpose-and-design-56ba6335778b>.


²⁸ The only limiting factor is the security of the distributed ledger, where scarcity depends on the robustness of the ledger's governance.

²⁹ Cryptoasset is a blanket term for any asset that is represented (i.e. issued and stored) in a distributed ledger. It is essentially a piece of data.

The most common execution architecture for a distributed ledger is a blockchain. Another type is the Directed Acyclic Graph, which we do not cover in this overview due to its currently limited use.³⁰

In the dimension of control, the sliding scale varies between completely centralised and completely decentralised. On the centralised end of the spectrum, there are private (permissioned) ledgers which are fully controlled and used by only one entity. On the decentralised end, there are public (permissionless) ledgers, where anyone can not only store data, but also help develop, manage and verify the ledger. Note that permissioned ledgers typically refer to those which are more open than private ones, but still require certain permissions to use or help maintain it.

Figure 6: Distributed Ledger Structures



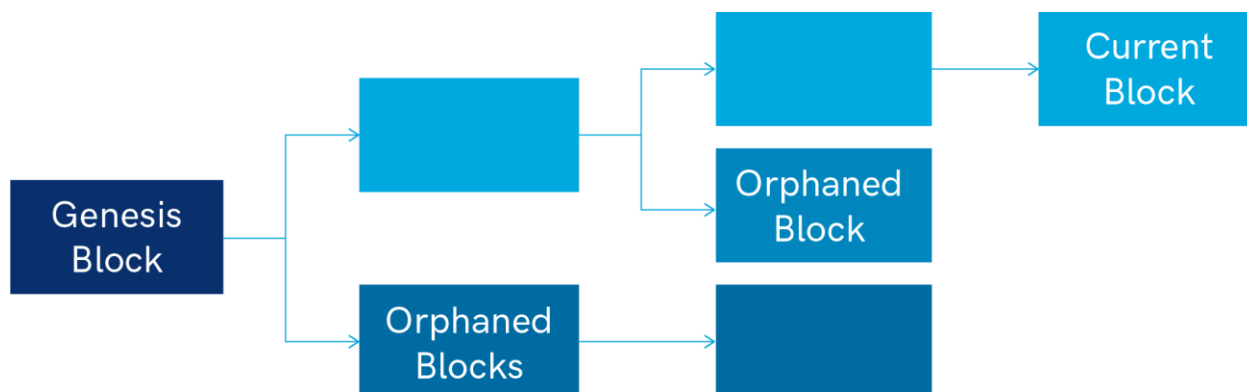
Source: Aaro Capital Research

2.4 Blockchains

The blockchain is the most common type of a distributed ledger. It consists of possibly several chains of blocks. Each block contains pieces of information, such as financial transactions. The order of blocks matters and the first is called the genesis block. Although several chains may coexist temporarily, there is consensus on the one that everyone follows and updates according to some rule defined by the blockchain protocol. The longest chain rule, a common consensus rule for blockchains, is shown in blue in Figure 7 below. Blocks which do not make it into the final blockchain are called orphaned blocks.

³⁰ The Directed Acyclic Graph solves some key issues faced by blockchain but has its own limitations. So far, it has received limited attention, however interesting projects using this data architecture include IOTA, NANO, Byteball and Hedera Hashgraph.

Figure 7: A blockchain with orphaned blocks



Source: Aaro Capital Research

When new transactions are generated, they are recorded in the blockchain. This is accomplished by a writer who creates a new block and attaches it to the last block of the consensus chain. If there are many potential writers, a rule specifies who is going to write the new block. For example, Bitcoin and Ethereum implement the Proof-of-Work protocol, covered in Section 5.1. Writers not only add new blocks but also maintain a copy of the ledger.

2.5 Public, Private and Permissioned Ledgers

As outlined in section 2.3 above, there is a sliding scale on the control dimension of a ledger. On this scale, ledgers typically fall into three buckets: private, permissioned and permissionless.

Private ledgers are controlled and used by only one entity, like an internal company database. In this setting, the advantages of a distributed ledger, such as robustness, immutably and transparency, can be achieved with traditional database structures which are far more efficient in terms of speed, resources and customisation.³¹ Therefore, there is little reason to use a private distributed ledger.

The term permissioned ledgers typically refers not to private ledgers, but to jointly controlled and maintained ledgers, with a controlled user base and a small number of semi-trusted nodes.³² This allows for greater ledger control, greater customisation and does not require a cryptocurrency or sybil resistance mechanism to align incentives.³³ Permissioned ledgers also use a different consensus algorithm than most permissionless ledgers. These ledgers are favoured by enterprises for business-to-business transactions.

Public, or permissionless ledgers, are open to everyone. Anyone can run a network node to verify their own copy of the ledger; they may choose to extend the ledger by competing in mining for blocks, as well as develop the open source code on which it runs. As anyone can attempt to extend the ledger, such contributions cannot be trusted to adhere to the rules - therefore, it is up to all users to verify each contribution against their own copy of the rules. Trust and faith must not be a requirement of the protocol as there is no accountability or governing

³¹ For more information, see: <https://www.multichain.com/blog/2016/03/blockchains-vs-centralized-databases/>.

³² Nodes are semi-trusted as a small number of known entities are preselected on the basis that they will act in the best interest of the ledger's users and administrator. However, byzantine fault tolerance consensus algorithms together with classical consensus algorithms used by permissioned blockchains only require $3n+1$ trustworthy nodes to guarantee non-compromised consensus, where n is the number of malicious nodes. For more information, see: <https://www.persistent.com/wp-content/uploads/2017/04/WP-Understanding-Blockchain-Consensus-Models.pdf>

³³ Sybil resistance mechanisms are covered in Section 5.

body. Permissionless ledgers use game theory to align the incentives of all users, including block generators, through various mechanisms, most notably rewards in a native cryptocurrency and penalties via external costs (e.g. expenditure of electrical energy in Proof-of-Work schemes). Accordingly, permissionless ledgers are known as trustless networks. There are tens of thousands of individual nodes on some of the larger networks. Due to their trustless nature and absence of large controlling entities, permissionless ledgers are typically used in business-to-customer or peer-to-peer transactions.

Table 5 below provides an overview of the advantages and disadvantages of permissioned and permissionless distributed ledgers. Currently, permissioned ledgers have clear advantages over their permissionless counterparts, and are therefore more readily adoptable by businesses in the short-to-medium term. These include higher throughput, faster and certain finality of transactions, and easier implementation of privacy. However, these advantages are achieved through increased centralisation, which somewhat undermines the *raison d'être* of distributed ledgers - to achieve trust between participants without giving away too much market power to trusted intermediaries - as there is still need to semi-trust nodes and platform administrators. Moreover, in the absence of a native cryptocurrency, monetisation of the platform can be tricky as it requires an entity to control the ledger, again undermining the trust benefits of distributed ledgers.

Permissionless ledgers are generally completely trustless and censorship-resistant by design. Furthermore, a native cryptocurrency can capture the positive externalities of network effects as ledger usage grows.³⁴ They therefore have the potential to better solve trust and market power issues across many use cases. However, while technical progress in terms of throughput, finality, privacy and sybil resistance mechanisms is promising, it will likely take several years before these shortcomings are adequately addressed.³⁵

In section 7, we discuss several trade-offs in the design of ledgers, along the dimensions of scalability, security and decentralisation. We conclude that permissioned and permissionless ledgers are both likely to continue to have widespread use, albeit in use-cases with different qualitative characteristics.

³⁴ The ability of a cryptocurrency to capture the positive externalities of network effects depends on the token design of the platform. Poorly designed token design will create a worthless cryptocurrency. This is expanded upon in section 5.

³⁵ The Libra project, a cryptocurrency stablecoin project being spun out of Facebook, will start with a permissioned ledger structure due to the current shortcoming of permissionless blockchains, but aims to ultimately become a fully permissionless ledger. For more information, see <https://developers.libra.org/docs/assets/papers/the-libra-blockchain.pdf>.

Table 5: Comparing control structures for distributed ledgers

	Private	Permissioned	Permissionless
Pros	Private	Higher throughput No inefficient sybil resistance mechanism More customisable and versatile Faster finality of transactions	Trustless Censorship resistant Removes market power of trusted intermediaries Efficient monetisation via the cryptocurrency No counterparty risk
Cons	Offers no benefits over traditional database structures which are quicker and resource efficient	Only semi-trustless Not censorship resistant Entities which retain control over the ledger may still exert market power Hard to monetise while still retaining the trust benefits of distributed ledgers Susceptible to human error by ledger admin Counterparty risk	Sybil resistance mechanisms can be inefficient Cryptocurrencies can be volatile Suspectable to attack by larger corporations or governments Less versatile or customisable Further development needed for: higher throughput, fast finality of transactions, regulation compliant privacy and user friendliness Finality of transaction is probabilistic (i.e. not 100% final)

Source: Aaro Capital Research, Investopedia³⁶, nakamo.to³⁷

³⁶ <https://www.investopedia.com/news/public-private-permissioned-blockchains-compared/>

³⁷ <https://medium.com/nakamo-to/whats-the-difference-between-a-public-and-a-private-blockchain-c08d6d1886a0>

3 Permissioned Distributed Ledgers

Although cryptocurrencies are the most well-known application of DLT, there are many others that may be easier to adopt in the short run. Note that many of the advantages that DLT offers over existing paper-based systems come from the digitisation and standardisation of data recording and data transfer, which is not unique to DLT. However, traditional shared databases have struggled to gain adoption due to trust issues between market participants, as discussed in section 1.

3.1 Distributed Ledger Platforms for Enterprises

Although permissionless DLTs can achieve transparency and decentralisation, enterprises often value privacy and control, together with the fast processing and finality of transactions. In recent years, there has been an emergence of several distributed ledger platforms for enterprises.

Hyperledger Fabric is a permissioned network that grants users with specific access rights.³⁸ It does not issue a currency, rather smart contracts called chaincodes. It achieves confidentiality by encrypting transactions, which can only be modified by authorised users. Another important feature is the modularity of the platform. Each project can use different components according to its needs, such as consensus and membership services. This means that fewer steps of verification are needed, thus minimising costs and optimising performance.

Corda is an open source DLT platform that allows businesses to transact directly with each other.³⁹ It was originally focused on financial enterprises; however, it now has a much broader reach. As with Hyperledger Fabric, it does away with the need for costly and time-consuming reconciliation in order to reach consensus. Moreover, it provides a framework for building applications called “CorDapps”.

Quorum is based on the Ethereum platform and was created through the introduction of the Enterprise Ethereum Alliance, a standards organisation with members such as Microsoft and JP Morgan.⁴⁰ Quorum aims at achieving transaction and contract privacy, together with the fast processing of transactions.⁴¹

3.2 International Remittances, Cross Border Transfers and Clearance of Payments

Using permissioned distributed ledgers, banks (or different subsidiaries of the same bank across countries) can overcome issues related to incompatible database systems, thus increasing efficiency and reducing costs. HSBC recently announced that in 2018 they cleared 3 million foreign-exchange transactions worth around \$250bn using DLT, which increased efficiency and speed and reduced their reliance on external technology providers.⁴² Ripple, which has also issued its own cryptocurrency, XRP, is working with banks in order to provide DLT solutions for instant clearance and settlement of payments. We provide more information on Ripple in section 4.3.

The international remittances market is worth at least \$600bn every year. However, it is also very fragmented. When a worker in the US sends money to their family in India, several intermediaries are involved, such as local

³⁸ For more information on Hyperledger Fabric, see <https://www.hyperledger.org/projects/fabric>.

³⁹ For more information on R3 Corda, see <https://www.r3.com/corda-platform/>.

⁴⁰ For more information on the Enterprise Ethereum Alliance, see <https://entethalliance.org/>.

⁴¹ For more information on Quorum, see <https://github.com/jpmorganchase/quorum/wiki/Quorum-Overview>.

⁴² More information can be found at: <https://www.ft.com/content/60d5a48c-17fa-11e9-9e64-d150b3105d21>.

banks in each country and a bank that handles the exchange between the two currencies. These intermediaries make the transaction slow and expensive.⁴³ Several startups are exploring the idea of using permissioned or permissionless ledger to clear transactions very quickly and considerably reduce transaction costs. Examples include BitPesa, Bitso and Circle.^{44,45,46} Libra, a project spearheaded by Facebook, plans to create a new stablecoin, called Libra, based on a decentralised DLT and smart contract platform. The aim is to allow Facebook's large user base to effortlessly send money to each other and make purchases on the platform.⁴⁷

A similar opportunity arises in the case of retail payments. At the moment, a merchant relies on intermediaries in order to confirm and clear payments. At the same time, one is responsible if the transaction turns out to be fraudulent. With DLT, however, there is no need for intermediaries, and it is extremely hard to corrupt a transaction or steal somebody else's identity.

Table 6: Payment clearing with and without DLT

Without DLT	With DLT
<ul style="list-style-type: none"> • Slow (up to three days) • Many intermediaries, each taking a cut • Transaction process is hard to trace and confirm • Transactions usually require manual intervention 	<ul style="list-style-type: none"> • Instant (within minutes or seconds) • One, low transaction fee • Transaction process is public and traceable • Transactions are automatic and can be programmed to execute conditional on certain events (e.g. via smart contracts)

Source: Aaro Capital Research

3.3 Trade Finance

Global trade finance transactions are worth \$10 trillion every year, yet most of the procedures are still paper based, resulting in a slow and inefficient process.⁴⁸ DLT has the potential to revolutionise this sector, initially by digitising all procedures. Moreover, the distributed ledger's immutability reduces the risk of fraud (e.g. in paper-based letters of credit) and speeds up the clearing of transactions. The transparency of information, and the fact that it is tamper-proof and cannot be altered, can motivate small and medium enterprises to share their information in the ledger. They also have full control of who accesses the information, thus avoiding the double counting of assets and transactions. DLT improves trust between participants via multiple points of verification. Finally, intermediaries for checking and verifying information would no longer be needed, thus reducing costs. All of these benefits can stimulate the access of these enterprises in world trade, thus boosting economic activity.

Banks have a large incentive to facilitate a boost of trade and an increase in trade finance transactions. There are at least two DLT platforms that facilitate the clearing of trade finance transactions. The We.Trade platform is

⁴³ More details can be found at <https://www.cbinsights.com/research/blockchain-disrupting-banking/>.

⁴⁴ For more information on BitPesa, see <https://www.bitpesa.co>.

⁴⁵ For more information on Bitso, see <https://bitso.com>.

⁴⁶ For more information on Circle, see <https://www.circle.com>.

⁴⁷ For more information, see the white paper at <https://developers.libra.org/docs/assets/papers/the-libra-blockchain.pdf>.

⁴⁸ See <https://www.tradefinanceglobal.com/trade-finance/> for an introduction.

a European digital trading platform, backed by Deutsche Bank, HSBC and UBS, among others.⁴⁹ The eTrade Connect platform is based in Hong Kong and developed by HSBC, Standard Chartered and ten other banks.⁵⁰

Table 7: Trade finance with and without DLT

Without DLT	With DLT
<ul style="list-style-type: none"> • Still paper-based • Slow process • High cost per transaction 	<ul style="list-style-type: none"> • Digital transactions, reduced risk of fraud • Instant clearing of transactions • Low cost, which can induce smaller companies to start exporting, thus increasing trade

Source: Aaro Capital Research

3.4 Supply Chains

DLT can contribute the most in industries where many different trading partners exist, each contributing their own information and controlling their own privacy, without any partner being particularly dominant. In many cases, the most relevant application would be a centralised, permissioned ledger.⁵¹ A prominent example is a distributed ledger that records transactions within a supply chain, thus providing efficient exchange of information between parties and, most importantly, increasing the reliability of data. Moreover, the distributed ledger allows each party to verify their identity, so that other parties can establish some trust before forming a business relationship.⁵²

The IBM Food Trust is a permissioned ledger which provides authorised users access to food supply chain data.⁵³ This includes the current location of a food item, where and when it was produced and under what conditions, what certifications it has obtained, and how it was transported. The fact that each contributing party can completely control who has access to the information it provides makes it difficult to manipulate or take advantage of the distributed ledger. This incentivises increased participation and facilitates a more efficient supply chain. For example, Walmart has used the IBM Food Trust as a tool to ensure food safety, by enabling users to trace food products through its supply chain.⁵⁴

Another example is the tracking and verification of the authenticity of luxury goods. The luxury brand conglomerate LVMH is developing a permissioned DLT, called AURA, in its fight against counterfeit goods. Using

⁴⁹ For more information on We.Trade, see <https://www.ibm.com/case-studies/wetrade-blockchain-fintech-trade-finance>.

⁵⁰ For more information eTrade Connect, see <https://www.etradeconnect.net/Portal>.

⁵¹ With a distributed ledger some meta data will inevitably be visible to other users of the platform.

⁵² An example is a lorry driver who bids for a contract to deliver some cargo. At the moment, it is very difficult to verify the driver's identity accurately, and it is possible for a lorry driver to falsify their identity and steal the cargo. This is currently a major issue in supply chains (see <https://losspreventionmedia.com/unreported-cargo-theft-incidents-make-it-difficult-to-grasp-scope/>). With a distributed ledger, the past transactions of the lorry driver are publicly visible, because they are recorded on the ledger as transactions of a specific public address. Using cryptography, only the real lorry driver can prove they own this public address and therefore has performed the previous assignments that led to the recorded transactions. This verification is instant and nearly impossible to falsify, unless someone else obtains their private key.

⁵³ More information at <https://www.ibm.com/downloads/cas/EX1MA1OX>.

⁵⁴ More information at https://mediacenter.ibm.com/media/Walmart%27s+food+safety+solution+using+IBM+Food+Trust+built+on+the+IBM+Blockchain+Platform/1_b3n798xc/98867192.

Quorum, a permissioned version of the Ethereum blockchain developed by JP Morgan, the DLT will provide Proof-of-Authenticity for luxury goods and trace their origins. The first two brands scheduled to participate are Louis Vuitton and Parfums Christian Dior. However, there are plans to include more brands and eventually competitor firms.⁵⁵

Table 8: Advantages of DLT for supply chains

Advantages of DLT
<ul style="list-style-type: none"> • Replaces siloed data with a distributed database, which facilitates the easy exchange of information between supply chain partners • Immutability of data, once entered, reduces the possibility of encountering faulty data • Identity enables users to identify and trust the source of data • Identity and technology (Internet of Things enabled sensors) allows direct oversight of different parts of the supply chain

Source: Aaro Capital Research

3.5 Insurance

Insurance is another industry where many participants are involved, each with their own private information and with privacy considerations, but without any single entity that maintains a centralised ledger. Although there are several regulatory and legal hurdles to overcome, there are many potential applications of DLT that can deliver significant efficiencies in the industry.⁵⁶

Insurance fraud (excluding health insurance) amounts to around \$40 billion per year in the US alone. One reason is that information is not efficiently shared between insurers, reinsurers and the insured. This creates several inefficiencies, such as enabling multiple claims for the same accident, falsely claiming ownership of assets through counterfeiting, or unlicensed brokers selling insurance. Porting all this information to a distributed ledger has the potential of saving money for insurers and reducing premiums for the insured. Etherisc is an example of a startup which sells flight delay insurance.⁵⁷ Payments are automatic after a qualifying event takes place, which is verified using oracles and smart contracts.

In the property and casualty insurance market, DLT can provide a faster and more efficient settlement of claims, by aggregating the information of multiple parties at the time of an accident. Insurewave is a DLT-powered marine hull insurance platform, backed by EY, Guardtime and A.P. Moller – Maersk, among others, which launched in 2018.⁵⁸

B3i Services is a startup aimed at exploring the use of DLT in the reinsurance industry and is backed by some of the biggest firms in the industry, such as AXA, Generali, Hannover Re and Allianz.⁵⁹ Their first product enables

⁵⁵ More information at <https://www.coindesk.com/louis-vuitton-owner-lvmh-is-launching-a-blockchain-to-track-luxury-goods>.

⁵⁶ More information at <https://www.cbinsights.com/research/blockchain-insurance-disruption/>.

⁵⁷ For more information on Etherisc, see <https://blog.etherisc.com/democratizing-insurance-using-blockchain-2cdac647e957>.

⁵⁸ More information about the property and casualty insurance market, including Insurewave, can be found at https://www.ey.com/en_gl/news/2018/05/world-s-first-blockchain-platform-for-marine-insurance-now-in-co.

⁵⁹ For more information on B3i, see <https://b3i.tech/home.html>.

the rewriting of reinsurance contracts as smart contracts on a distributed ledger.⁶⁰ When an event occurs, such as an earthquake or hurricane, it is independently verified by an oracle and then the smart contract executes automatically, allocating payments across parties.

Table 9: Insurance markets with and without DLT

Without DLT	With DLT
<ul style="list-style-type: none"> Insurance fraud by placing multiple claims is common, because different pieces of information are stored in siloed databases Verifying whether an event occurred is difficult and expensive Processing a claim is slow and requires manual input Insurers, reinsurers and the insured use siloed databases which are not interoperable, resulting in duplication of data processing and storage 	<ul style="list-style-type: none"> All data is stored on the distributed ledger and each party has specific rights on which parts it can read and write, each having direct access to the data they need to perform their service There is the potential of incentivising independent oracles to verify events that trigger clauses in a contract Processing a claim can be done quickly and automatically, using a smart contract and independent oracles

Source: Aaro Capital Research

3.6 Healthcare

As outlined in section 1, distributed ledgers have several key economic advantages over previous attempts to create seamless shared databases, such as the one envisaged by the US HITECH Act in 2009. The healthcare industry is ripe for disruption, as DLT can offer several advantages to participants that are currently not available.⁶¹ Most importantly, DLT can allow a user to own their own health data and decide which interested party has access to it and when, for example a doctor or an insurance firm. Since a distributed ledger is append-only, health data can be tamper-proof, increasing verifiability and value. There is greater scope for consistency, as the data can exist as a single entry in a distributed ledger, rather than different versions in various siloed databases. Greater transparency about how data is used can lead to more individuals sharing their own data, thus increasing the value of the distributed ledger. Furthermore, it removes the duplication of often manually inputted data across private databases.

The integration of DLT in healthcare is proving to be a slow process as there are many hurdles to overcome, both regulatory and in terms of convincing interested parties to share their data. However, there are several early attempts that are worth noting. Guardtime is a start-up which has started implementing electronic health records using DLT.⁶² HealthVerity provides a health data exchange marketplace, using DLT to manage permissions and access rights.⁶³

⁶⁰ For more information, see <https://b3i.tech/what-we-do.html>.

⁶¹ A longer report is available at <https://www.cbinsights.com/research/report/blockchain-technology-healthcare-disruption/>.

⁶² For more information on Guardtime, see <https://guardtime.com/health>.

⁶³ For more information on HeathVerity, see <https://healthverity.com/>.

Table 10: Healthcare markets with and without DLT

Without DLT	With DLT
<ul style="list-style-type: none"> • The medical records of a patient are stored in private data silos which are difficult to access and share • A patient needs to seek permission to access their own medical records • Databases of medical records across institutions are not compatible, difficult to aggregate and may not be consistent • Opaque rules about access of medical records disincentivises participation and sharing • Duplication of data entry and processing for each individual database 	<ul style="list-style-type: none"> • Each patient has complete control over their own medical records, which are encrypted and stored on the blockchain • A patient grants or revokes access to their own data • Greater transparency and control over medical records can increase participation and sharing • Removal of duplicate data increases consistency and reliability • Time stamps, immutability and cryptographic identity of blockchains increase the verifiability of patent records

Source: Aaro Capital Research

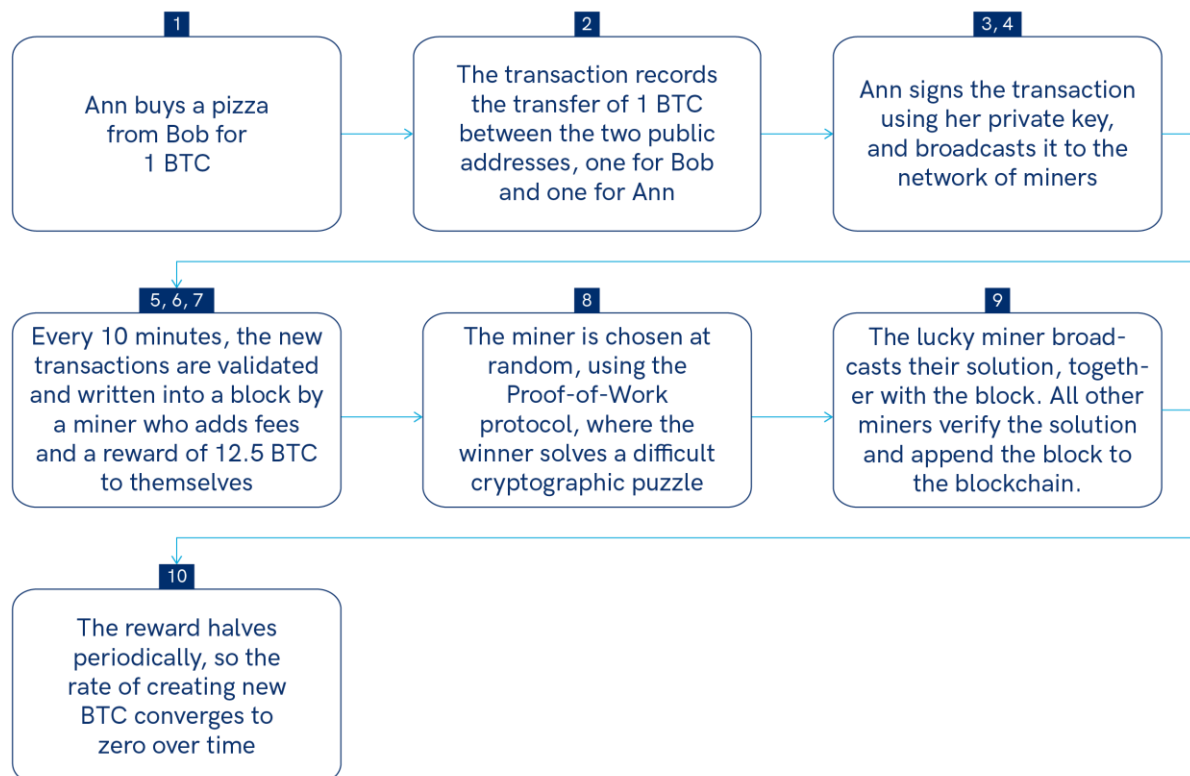
4 Permissionless Distributed Ledgers

The aim of permissionless distributed ledgers is to dispense entirely the need for trusted third party intermediaries and market failures that they create. This is accomplished by issuing cryptocurrencies and using tools from game theory and cryptography to incentivise participants to act in the best interest of the network. We review four prevalent permissionless ledgers and their cryptocurrencies, each with distinct design characteristics - Bitcoin, Ether, XRP and Zcash.

4.1 Bitcoin

Bitcoin (BTC) is the first, most well-known and largest cryptocurrency, implementing a permissionless and distributed blockchain. The cryptocurrency's primary function on top of the bitcoin ledger is to act as an incentive and coordination mechanism that prevents attacks that corrupt the data stored in the ledger.⁶⁴ There are multiple copies of the distributed ledger, maintained by different pseudonymous participants. The ledger consists of a list of all transactions (although other types of data can also be added), and in principle anyone is permitted to write and read information on the ledger and verify that the transactions are genuine. We explain how the mechanism works, using an example of a transaction between Ann and Bob, outlined in Figure 8 below. The numbers of each stage correspond to the numbering of the more technical explanation provided in Figure 9. However, Figure 9 is not intended to cover all technical intricacies of this process.

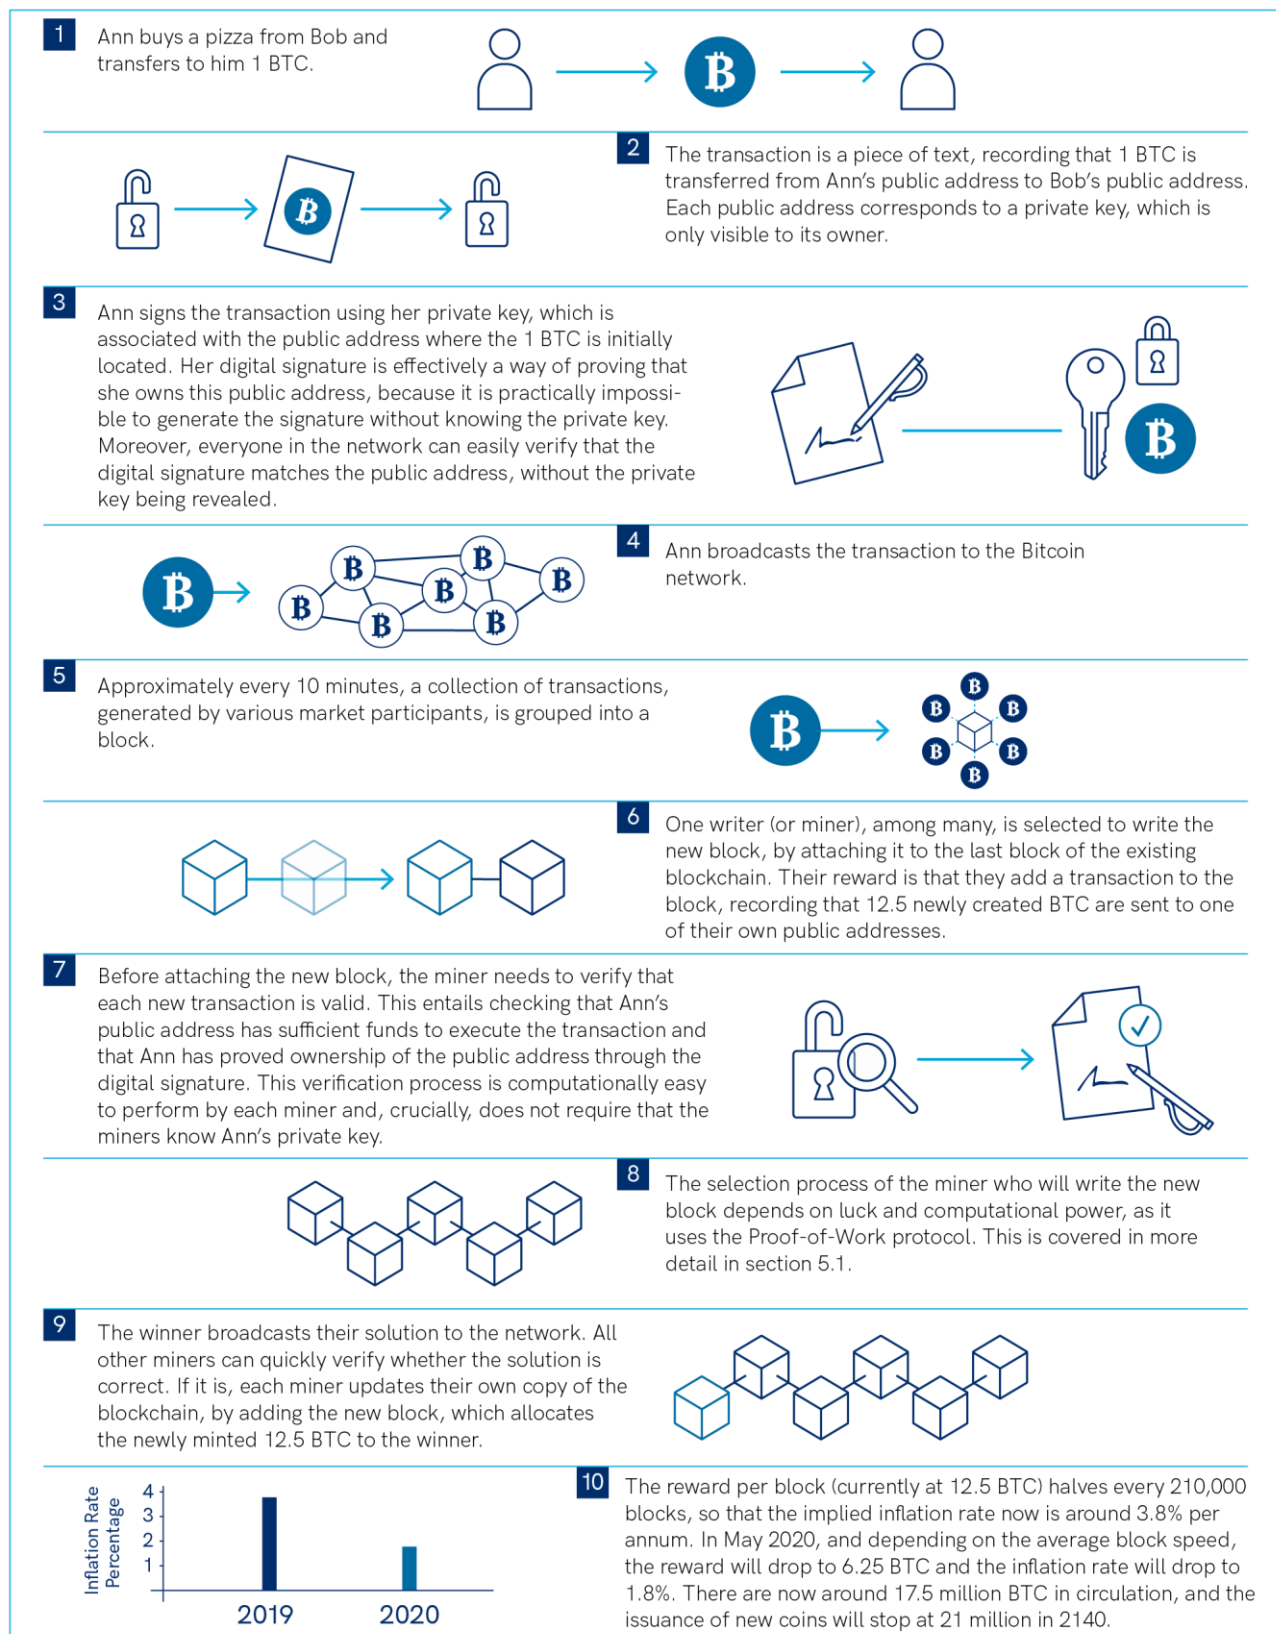
Figure 8: Simple overview of the bitcoin transaction process



Source: Aaro Capital Research

⁶⁴ The most well-known type is the 51% attack, which we explain in Section 5.4.

Figure 9: Technical overview of the bitcoin transaction process



Source: Aaro Capital Research

Miners also charge transaction fees on top of the block reward.⁶⁵ If Ann wants to transfer 1 BTC to Bob, she needs to specify a higher amount, for example 1.01 BTC. The remaining 0.01 BTC residual is the transaction fee to the miner who will write this new block. The fee is chosen by Ann. However, it is up to the discretion of the miner to whether or not to confirm a transaction with a low fee.

It is important to note that every BTC in circulation is created initially as a reward to a miner, who successfully solves the cryptographic puzzle first and attaches the new block of transactions to the consensus blockchain. However, these BTC “exist” only as long as this new block is part of the consensus blockchain in the future. This creates the incentive for every miner, who has ever been rewarded for mining and still holds BTC, to defend this consensus blockchain against malicious participants who might want to create a fork - an alternative branch of the blockchain. As the blockchain grows longer, the computational (thus economic) cost of corrupting it becomes higher.

A potential issue with Bitcoin is the relative concentration of miners. In 2014, one miner controlled close to 50% of the computing power (or hash rate), thus making the network vulnerable to attacks.⁶⁶ In 2019, however, no miner controls more than 20% of the hash rate, leading to a Herfindahl-Hirschman Index of around 1200, classifying the market as competitive.^{67,68} There has also been a high turnover of bitcoin mining pools over this period, further suggesting that the market is currently competitive.⁶⁹

4.2 Ethereum

The Ethereum blockchain and its associated cryptocurrency, Ether, is similar to Bitcoin, in that it implements public blockchain technology to verify transactions and maintain a distributed ledger. Moreover, it currently uses the Proof-Work protocol. However, Ethereum is different because it provides an additional layer of infrastructure, a virtual machine, which enables developers to embed complex logic in the form of smart contracts on the blockchain.⁷⁰ These smart contracts are executed in a trustless manner by all network participants. Storing some data and logic on a public blockchain is what differentiates so-called “dApps” from more familiar web apps. This is illustrated in Figure 10 below. Whereas Bitcoin’s main purpose is to be a universal means of payment and store of value, Ethereum’s purpose is to be the world’s distributed computer. Programmers can concentrate on building dApps for a variety of uses, on top of an infrastructure that has solved the issues of consensus, mining, storage and computation. To put this into perspective, traditional app developers are building apps without worrying about the issues of scalability, storage and computation, but with the added caveat that their app is hosted by a

⁶⁵ Fees will become more important as fewer new coins are created and the bitcoin inflation rate converges to 0. However, fees are still required today to incentivise miners not to mine empty blocks. A more technical discussion on how mining fees work can be found here: <https://hackernoon.com/blockchain-fees-are-broken-here-are-3-proposals-to-fix-them-1f772e1530dd>

⁶⁶ There are many types of attacks. The most prominent one is a 51% attack, which we explain in Section **Error! Reference source not found.**

⁶⁷ For more analysis on Bitcoin mining HHIs, see: <https://ark-invest.com/research/ark-disrupt-issue-144>.

⁶⁸ For more information on the Herfindahl-Hirschman Index, see <https://www.investopedia.com/terms/h/hhi.asp>.

⁶⁹ More information on miner turnover can be found at: <https://a16z.com/2019/02/09/voting-blockchains-governance-security-cryptoeconomics/>.

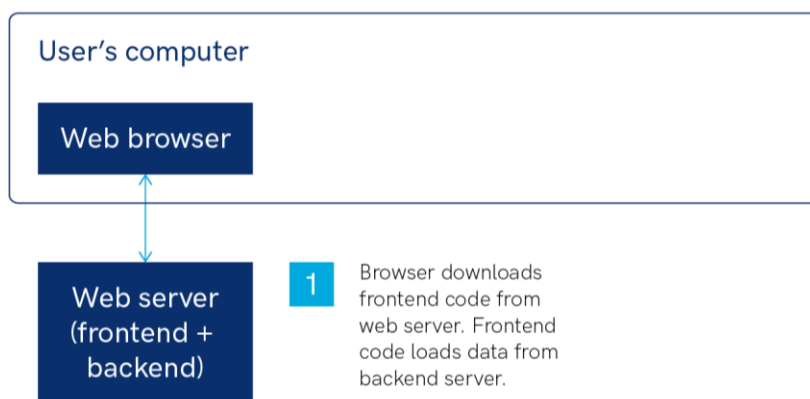
⁷⁰ Bitcoin and other distributed ledger protocols have their own virtual machines which allow for limited smart contracts e.g. exchange of on-chain assets and cryptographic key validation. Smart contract platforms such as Ethereum offer Turing Complete virtual machines where, in the case of Ethereum, gas serves as the limitation to the number of transactions which can be performed.

centralised provider, such as Amazon's AWS or Microsoft's Azure Cloud. The promise of Ethereum is to provide these services in a trustless and distributed environment, essentially creating a Web 3.0 platform.⁷¹

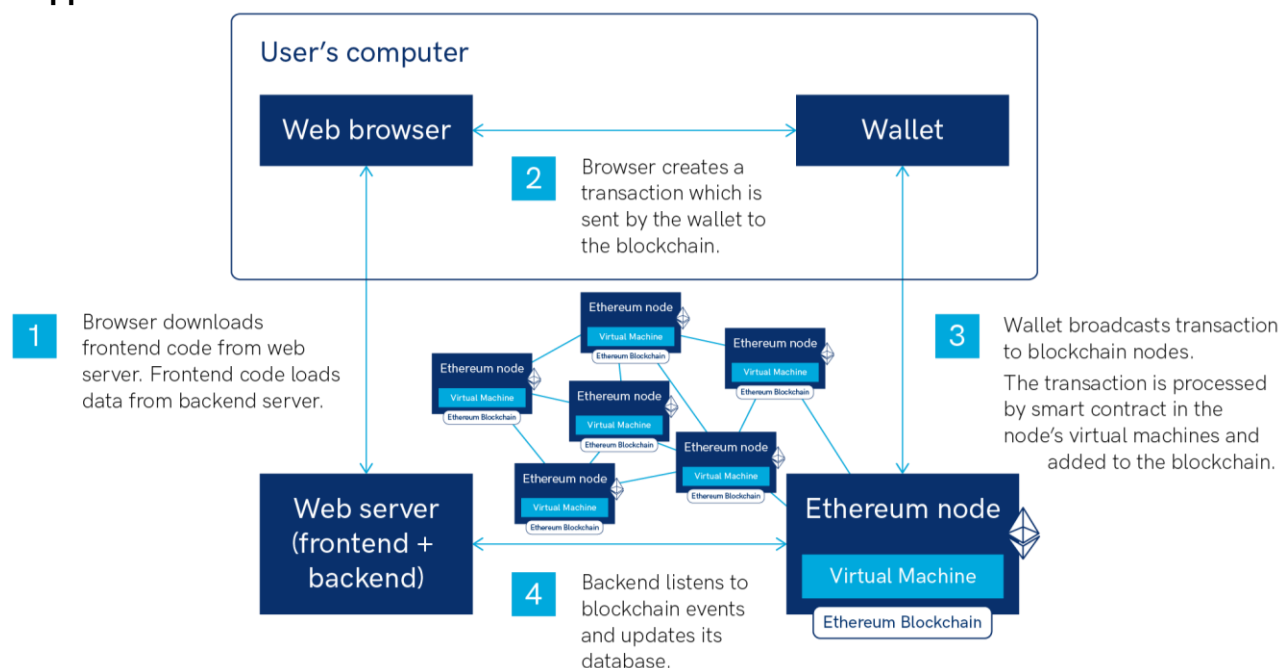
Smart contracts are programs that are coded on the blockchain. They provide a set of conditions, recorded on the blockchain, which trigger automated actions when satisfied. Decentralised applications use smart contracts and the computing power of the network in order to perform some functions. To provide an analogy, we can think of the Bitcoin blockchain as a dApp whose function is to record financial transactions, using a token called Bitcoin. Beyond this, there may be a multitude of functions that can be implemented in a decentralised and trustless environment.

Figure 10: Web based apps vs decentralised apps

Web Based Apps



dApps



Source: Aaro Capital Research

⁷¹ This is covered in more detail in the paper "An introduction to Web 3.0".

Case Study: Augur

Augur is a set of smart contracts on the Ethereum blockchain. A user of Augur can create a prediction market, whose aim is to forecast future events, by leveraging the wisdom of the crowd. This is based on the premise that a large number of people will collectively have more information about the probability of an event than a small number of experts. Thus, a prediction market is a mechanism that incentivises participants to aggregate their private information in order to collectively form a probability. For instance, suppose that we want to get an estimate of the probability that the earnings of a company in the next quarter will increase or decrease. We create a prediction market with an asset that pays 1 ETH if they increase and 0 ETH otherwise. The initial price of the asset is 0.5 ETH, interpreted as the probability of the earnings increasing. Whenever a participant in that market thinks an increase is more probable than the current price, they buy the asset, otherwise they sell it. When the event occurs, a trusted source (usually called an Oracle) informs the market whether earnings increased or not, so that the asset pays accordingly. Participants buy and sell the assets using ETH. More importantly, the buy and sell orders are recorded in a smart contract, which also ensures that ETH is paid out when the event occurs, and information is revealed. The computations that allow this smart contract to operate are performed by nodes in the Ethereum network, which are also rewarded with ETH.

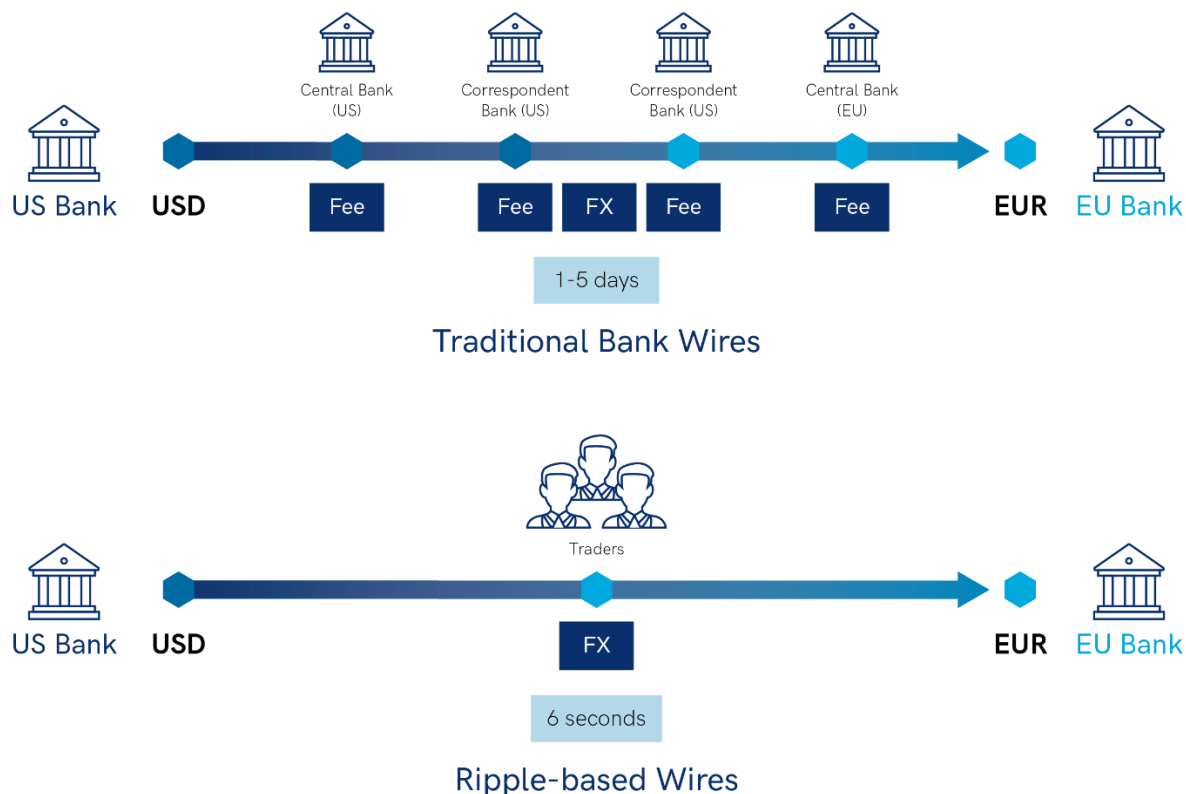
4.3 Ripple

Ripple is a company that provides a digital payment system that is being tested by several financial institutions and has recently been successful in some commercial applications. The system is based on the cryptocurrency XRP, which was created by Ripple. It has several differences relative to Bitcoin. First, although XRP also has a fixed supply (100 billion), it was all created at inception, and Ripple owns most of them. Hence, there is no mining, but transactions are still recorded in a blockchain. Instead of using the Proof-of-Work protocol, it uses a low-latency Byzantine agreement protocol, which can reach consensus without full agreement of all nodes.⁷² As a result, transactions settle very quickly within 4 seconds, as compared to 1 hour for Bitcoin (for 6 blocks to be generated), and more than 2 minutes for Ethereum.⁷³ Moreover, XRP is more scalable, as it can currently handle around 1500 transactions per second, as compared to 6-7 for Bitcoin. The intended use of XRP is as a bridge currency that facilitates foreign exchange and business-to-business payments.

⁷² A technical analysis can be found at <https://arxiv.org/abs/1802.07242>.

⁷³ <https://ripple.com/xrp/>

Figure 11: Traditional Bank Wires vs Ripple-based Wires



Source: Ripple

Ripple's blockchain, called RippleNet, involves a network of more than 200 banks and payment providers. It contains three main services: xCurrent (payment processing system for banks), xRapid (facilitates fast currency exchange using XRP) and xVia (facilitates business-to-business payments). There is now a small but increasing number of financial institutions using xRapid (and XRP) to complete commercial payments across countries.⁷⁴

4.4 Zcash

Zcash is a cryptocurrency focused on the privacy of transactions. It uses the Proof-of-Work protocol, just like Bitcoin, however transactions are recorded differently on the blockchain. Bitcoin transactions are always between two or more public addresses, so it is straightforward to trace the journey of each BTC, even though the owners of the public address may not be revealed.⁷⁵ To provide an analogy, it is as if all USD transactions between bank accounts are publicly announced, even though the owner of each account is not. The design of Zcash ensures that transactions between accounts can be made private.

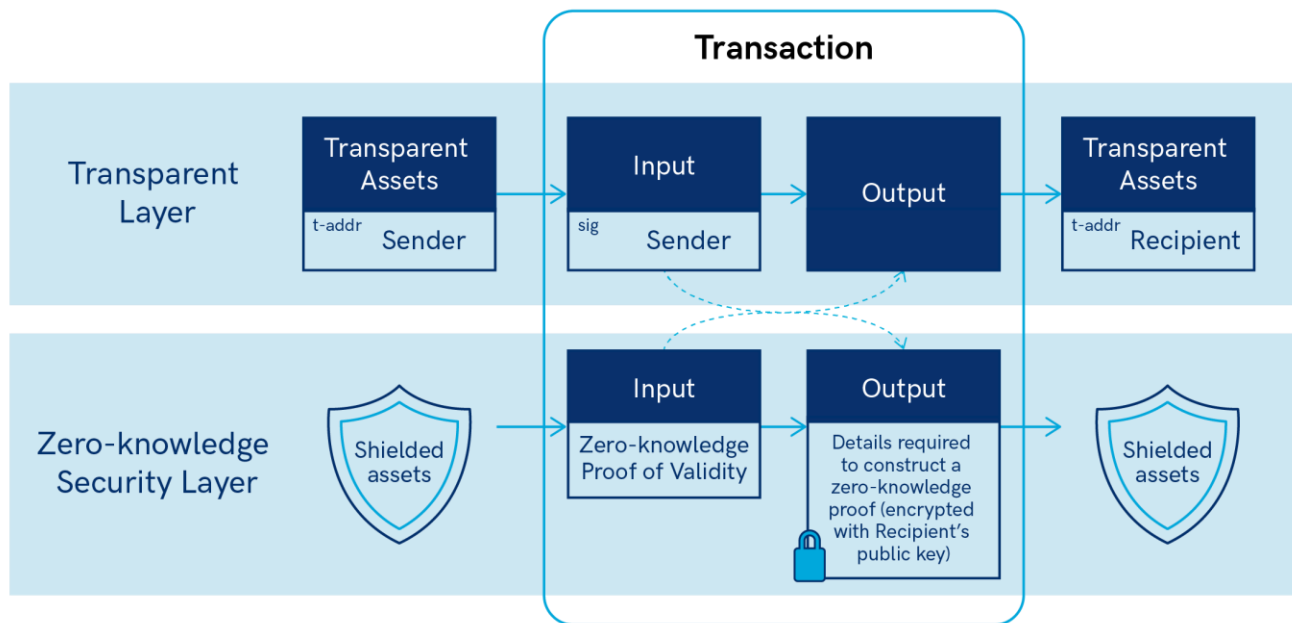
To achieve this, Zcash uses zero-knowledge proofs and two types of addresses: private (z-addresses) and transparent (t-addresses), where the latter is similar to the public addresses of Bitcoin. A transaction can be Z-to-Z, meaning that it is recorded on the public blockchain and known to have occurred, however the amount, the

⁷⁴ For more information on Ripple's partnerships, see <https://decrypt.co/5313/complete-ripple-partnerships-xrapid-xrp>.

⁷⁵ A caveat to this is a bitcoin mixer, which mixes coins between different addresses so that it is not possible to determine the exact senders and receivers.

fees and the addresses are encrypted and private. A T-to-T transaction is similar to a transaction recorded in Bitcoin, where the addresses, the fees and the amount are public. Moreover, Z-to-T and T-to-Z transactions are also possible.

Figure 12: Zcash's layered architecture



Source: Electric Coin Co

Although Zcash allows a transaction to be private, it is possible for the parties involved to provide some information for audit or compliance purposes. The aim of Zcash is not to facilitate illegal behaviour but to protect privacy as a fundamental right.⁷⁶

⁷⁶ In principle, privacy coins cannot ensure that no information will leak, because it still may be possible to infer details about a specific Z-to-Z transaction by combining public information about related T-to-T, Z-to-T and T-to-Z transactions. Moreover, as with all blockchains, exchanging fiat money to buy a privacy coin can be traced.

5 Consensus

One of the most important issues in DLT is how consensus, on the correct state of the ledger, is achieved among the many participants who maintain and update it. Since participants do not know the identity of others, how can they communicate and agree on what information is to be written and by whom?

The easiest way of choosing the writer of the next block is to randomly pick one participant.⁷⁷ However, this opens the possibility of a “Sybil attack”, where a participant creates multiple selves (e.g. multiple IP addresses) in order to increase their probability of selection and the payoff that they will receive. If a participant greatly increases their probability of selection, they can control the ledger for their own benefit and to the detriment of everyone else. The Proof-of-Work and Proof-of-Stake protocols are solutions to this problem and are therefore called “Sybil resistance mechanisms”. They generate scarcity of resources, making it increasingly difficult and expensive for a participant to create multiple selves. The Proof-of-Work protocol achieves resistance by selecting the participant (miner) who can first solve a difficult (costly in terms of computation) problem. Proof-of-Stake specifies that the probability of selection is proportional to the miner’s stake of coins, which are by construction scarce and cannot be replicated.

The other main issue is reaching consensus on which branch of the blockchain the new block of information is going to be attached. If there are two competing branches, how can the participants agree on which is the correct one? The two branches might have been created because of lack of communication and latency, or because some malicious participants altered the information in previous blocks in order to make their branch the consensus one. There are several ways of resolving this problem. Most permissionless blockchains use Nakamoto Consensus, or the longest chain rule, while permissioned blockchains use classical consensus.

5.1 Proof-of-Work

The Proof-of-Work protocol achieves resistance to Sybil attacks by selecting the participant (miner) who can first solve a difficult (and costly in terms of computation) problem. Each miner attempts to be the first to solve a difficult cryptographic puzzle. Its solution is a number which, when combined with the text of the previous and new block of transactions, produces a “hash”, beginning with a predefined number of zeros. There is no analytical solution to this problem, such that the only way of solving it is by trying many different combinations of numbers. The more (and quicker) computers a miner has at their disposal, the higher the probability is that they will find an acceptable solution first.⁷⁸ Crucially, when a miner finds a solution, they create a new block of transactions allocating the reward (newly minted coins and miner fees) to themselves, and broadcast the new block along with their solution to the wider network of miners. Other miners can instantly verify the solution, and thus verify if the miner has done the required work and incurred the associated cost to solve the problem.

Bitcoin and Ethereum currently use Proof-of-Work to select the miner who has the right to mine the next block and be rewarded with newly minted coins and miner fees. These coins provide incentives so that miners act in the best interest of the platform’s users. If the cryptocurrency platform is successful, its price will be high and each miner will be able to afford specialized equipment which has no use outside of mining, together with the

⁷⁷ In contrast, Byzantine Fault Tolerance algorithms require far more channels of communications between participants. This quickly limits the number of participants that can be practically involved in the decision-making process.

⁷⁸ One has to find a solution within a given margin of error of the exact solution.

associated electricity costs.⁷⁹ If the platform falls in popularity due to malicious miners, the value of the cryptocurrency will fall and miners stand to make a loss given the costs incurred to mine the cryptocurrency.⁸⁰

High electricity consumption has been one of the main criticisms of Proof-of-Work and Bitcoin in general.⁸¹ However, to put the cost into perspective, one needs to also consider the benefits. The most important benefit is creating a system that has never been corrupted up to now, where information and wealth can be stored without the use of trusted intermediaries.⁸² Moreover, the difficulty of finding a solution to the problem can increase or decrease by adjusting the allowable margin of error from the exact solution. In practice, for Bitcoin the difficulty is adjusted every two weeks, such that each problem takes on average 10 minutes to be solved. If there are too many orphaned blocks, meaning that two or more block producers were able to solve the problem almost simultaneously, the margin of error is reduced and the problem becomes more difficult. If a solution takes on average more than 10 minutes to be found, then the difficulty is reduced. This means that the cost of mining depends on the price of BTC, rather than the other way around. If many market participants use Bitcoin and its price is high, then there will be many miners competing to find a solution, hence the difficulty will increase and the energy cost will be high.

Finally, a recent report dispels the myth that Bitcoin mining has a large, detrimental environmental impact.⁸³ It finds that Bitcoin mining is powered on at least 74% renewable energy, such as solar, wind and hydro power. In particular in China, where a significant share of Bitcoin mining takes place, there is excess capacity in renewables that would otherwise be wasted.⁸⁴

5.2 Proof-of-Stake

The most prominent alternative to Proof-of-Work is the Proof-of-Stake protocol, currently used by cryptocurrency EOS (in delegated form) and planned to be adopted by Ethereum.⁸⁵ Proof-of-Stake specifies that the probability of selection is proportional to the miner's stake of coins, which are by construction scarce and cannot be replicated. Effectively, the blockchain first records a set of validators. A validator can be anyone who locks their

⁷⁹ If it is not sufficiently costly to perform a Sybil attack, it may be in the miner's best interest to double spend when they control 51% or more of the network's mining power. There may also be external motives to destroy the blockchain, for example in order to short the cryptocurrency.

⁸⁰ It is possible for miners to use their computer equipment to mine other cryptocurrencies more profitably. However, over time the mining hardware becomes more specialized to each Proof-of-Work algorithm, making it harder to mine different cryptocurrencies.

⁸¹ Another criticism of Proof-of-Work is that miners should engage in an activity that is socially useful, rather than conducting pointless mathematical calculations. Such an alternative Sybil resistance mechanism is the Proof-of-Space protocol, where the miner is selected based on their hard drive storage. Although not widely tested yet, it is less secure than Proof-of-Work, because the unpredictable evolution of the hard drive market may have an impact on mining. Moreover, it is cheaper to coordinate a 51% attack, because the attackers can always resell their hard drives after the attack, hence decreasing the cost of such an attack.

⁸² The benefits of this for users are covered in "*An introduction to Crypto's near Money Characteristics*" and "*An introduction to Web 3.0*".

⁸³ The report can be found at <https://coinshares.co.uk/wp-content/uploads/2019/06/MiningWhitepaperJun2019FinalForeword.pdf>.

⁸⁴ An earlier report at <https://coinshares.co.uk/wp-content/uploads/2018/11/Mining-Whitepaper-Final.pdf> finds that in solar and wind there is significant and fluctuating excess capacity in several Chinese provinces. For example, in Gansu the share of electricity produced by solar and rejected by the grid was 30% in 2015 and 2016, 20% in 2017 and 11% in 2018.

⁸⁵ More details about the adoption by Ethereum are provided at <https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQ>.

coins in a deposit. To provide an analogy, with Proof-of-Work, miners commit their computational power, whereas with Proof-of-Stake, validators commit their coins. As a result, staked coins are assets that can yield interest.

There are several proposed methods of selecting who is entitled to write the new block among the set of all validators. In chain-based Proof-of-Stake, the protocol randomly selects the winner, usually with the probability being proportional to the size of each validator's stake. In Practical Byzantine Fault Tolerance style Proof-of-Stake, a validator is randomly selected and given the right to propose the next block.⁸⁶ All other validators then vote and if there is a majority (or a super majority of at least two thirds), the block is accepted to be attached in the blockchain.

There are several potential benefits of Proof-of-Stake. Most importantly, there is no longer the need to incur electricity and computational cost in order to maintain and expand the blockchain, as in Proof-of-Work. Market participants are incentivised to stake their coins and not sell them in order to receive interest, which contributes to the stability of the network. This could create inequality, however, as those who initially hold the majority of coins will earn the highest share of new coins. Also, miners in Proof-of-Work are not incentivised to hold coins in the long term, but to invest in buying computational power.⁸⁷

The cost of coordinating and sustaining an attack in Proof-of-Stake is lower than in Proof-of-Work, as there is no cost of electricity. This could lead to more frequent attacks. What mitigates this problem is that it is easier to impose penalties to malicious participants after a failed attack, for example by confiscating their coins. In Proof-of-Work, this would amount to confiscating computers, which is impossible. However, if an attacker succeeds in substantially decreasing the price of a cryptocurrency, the value of all their coins is also diminished, while in Proof-of-Work an attacker has mining hardware which can be used to attack another cryptocurrency. In practice, Proof-of-Stake does not have a long track record of resisting attacks, as it has not been widely adopted yet.

Finally, Proof-of-Stake may suffer from the Nothing-at-Stake problem. This problem specifies that, since it is free to create new blocks, there is an incentive for a miner to participate in many different branches of the blockchain, even those that are corrupted, as long as there is even a small probability that one of them will be the consensus one. Such behaviour can create multiple branches and weaken the enforcement of consensus on a unique branch of the blockchain. This problem can be mitigated by imposing penalties on stakes which approve blocks that do not eventually get accepted into the consensus chain.

5.3 Nakamoto Consensus

A major issue in the design of a ledger is how participants reach consensus on which branch of the blockchain they will write the next block. This is important because if a branch becomes orphaned and no more blocks are added to it, all coins recorded on it are essentially worthless.

What is the consensus mechanism for agreeing on the "correct" branch of the blockchain? In principle, malicious participants could try to pass a corrupted block as the accepted one and convince everyone else that this is the case by providing false information. A system that avoids this problem satisfies the property of Byzantine Fault Tolerance.

⁸⁶ More details are provided at <https://medium.com/tendermint/a-to-z-of-blockchain-consensus-81e2406af5a3>.

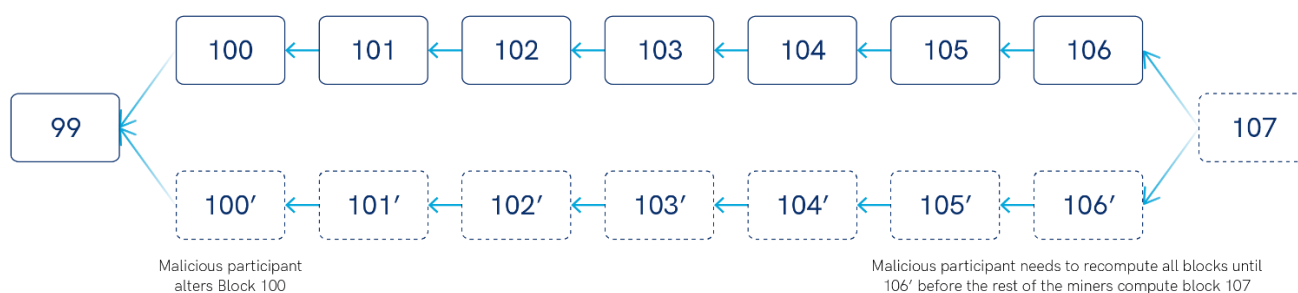
⁸⁷ The lack of direct benefit of holding coins beyond the need to pay minor fees gives rise to the potential for the velocity problem. Quantity theory of money states: $\text{price} = (\text{money supply} \times \text{velocity}) / \text{transactions in the economy}$. Without any reason to hold cryptocurrency long-term, there is little relationship between cryptocurrency usage, as highlighted at <https://multicoin.capital/2017/12/08/understanding-token-velocity/>.

There are two main solutions and several variations. The Nakamoto consensus specifies that agreement should be on the block that has the highest Proof-of-Work on it. This is also called the longest chain rule. In practice, it is agreed that miners always adopt the block with the highest number, because usually these are the blocks that have the most Proof-of-Work on them. Given that everyone else follows this rule, it is a best response for a miner to do the same. In other words, the Nakamoto consensus is a Nash equilibrium.

To see how the longest chain rule can hinder attacks, consider the double-spending problem. A malicious participant sends 1 BTC to buy some goods. Their aim is that, after receiving the goods, they rush to falsify the block that includes the transaction, to make it appear as if they never spent that 1 BTC. Suppose that the initial transaction occurs in block 100 and that the seller waits until block 106 to send the goods, such that 1 hour has elapsed (because each block is written every 10 minutes).⁸⁸

The malicious participant needs to alter block 100, erasing their transaction. In order to mine the alternative block 100', they still need to solve the cryptographic puzzle and consume computer power and electricity. As explained in section 4.1, the solution to the cryptographic puzzle depends on the text of both the previous and new block. This means that they need to solve again a new cryptographic puzzle for block 101', taking as input the text of the alternative block 100', thus expensing even more electricity. This process has to continue until they recompute block 106'. More importantly, they have to recompute all blocks until 106', before the other miners compute block 107, taking 106 as given. Because the cryptographic puzzle is very hard, it is almost impossible for a minority of miners to compute six blocks faster than all other honest miners can compute one block. In that sense, the combination of the Proof-of-Work protocol and the longest chain rule enforces consensus on the correct branch of the blockchain. To date, no such double-spending attack has been successful in the Bitcoin blockchain.

Figure 13: Altering a transaction



Source: Aaro Capital Research

A double-spending attack could occur if miners holding the majority of the computing power collude to falsify the blockchain. Since they have the majority of the computing power, they might be able to recompute several blocks before the honest minority computes the next block. This is called a 51% attack. In that case, however, the price of the BTC will drop and the value of the miners' wealth will diminish significantly. Since the majority of miners

⁸⁸ This "escrow" period of 6 blocks, before accepting that a transaction cannot be reversed, is usually followed in practice. This means that a transaction is never deterministically final, but with probability that converges to 1. However, most users consider that after 6 blocks, the probability of an additional block being rejected is sufficiently small enough to be able to accept it as final.

have already mined most of the previous blocks, have incurred the high entry costs of buying the mining equipment and have pocketed the block rewards, they have little incentive as a group to attack the blockchain.

The main deterrent against a 51% attack is the value of the cryptocurrency. The higher the value of the coins that miners currently hold, and expect to hold in the future by maintaining the blockchain, the greater the incentive to spend on computing resources. Thus, an increasing expense is required on hardware to maintain a controlling 51% of the network. In other words, the more valuable the cryptocurrency is, the more secure it should become. There have been relatively few successful 51% attacks on lower value blockchains, but a recent example is on Ethereum Classic in 2019.⁸⁹

An alternative to the longest chain rule, called classical consensus or Practical-Byzantine-Fault-Tolerance, follows the procedure below:⁹⁰

1. One player is randomly picked (for example in a Proof-of-Stake protocol) to propose a block.
2. All other players vote on whether to accept this proposal.
3. If the majority (or supermajority of 2/3rd of votes) accepts, then the block is accepted as the correct one.
4. Once a block is accepted, there is no reversal.
5. If there is no majority, then the system halts and there is a new proposal.

5.4 Security of Permissionless Ledgers and Traditional Databases

Permissionless ledgers predominantly implement the Proof-of-Work protocol together with the Nakamoto consensus, or longest chain rule. How secure is this system, as compared to the security of a centralised, permissioned database, for example that of a bank? Theoretically, a system can be breached in one of three ways. First, there is a human error by one or more insiders, who are entrusted with maintaining the system.⁹¹ This error could be unintended or incentivised by a malicious outsider. Second, there is a software vulnerability or insufficient security protocols that outsiders discover and exploit. Third, there is a brute force method of hacking the system, despite the efforts of the insiders.

In practice, databases are compromised very often. The overwhelming majority of these breaches are due to either a human error, or software vulnerabilities and weak security protocols. Both issues arise because of human mistakes.

The innovation of the permissionless ledger is that it removes the human factor as much as possible. There is no insider that alone maintains the system, hence their actions are irrelevant to security. Software vulnerabilities could compromise the system but, since the code is usually open source, mistakes can be checked by anyone and identified quickly. Moreover, there is a strong incentive for network participants to resolve any security issue in a responsible way, in order to protect the value of the coins they own. The only other method of hacking the ledger is by using the brute force of a 51% attack. The Proof-of-Work however, makes the cost of such an attack proportional to the value stored in the blockchain. In the case of Bitcoin, if the price is high and several miners


⁸⁹ The incident is described at <https://qz.com/1516994/ethereum-classic-got-hit-by-a-51-attack/>. Bitcoin Gold (which is different from Bitcoin) was attacked in 2018, as explained at <https://qz.com/1287701/bitcoin-golds-51-attack-is-every-cryptocurrencys-nightmare-scenario/>. The paper at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3290016 documents attacks on 13 coins.

⁹⁰ Several variations of this procedure have been proposed. A more detailed analysis can be found at <https://medium.com/tendermint/a-to-z-of-blockchain-consensus-81e2406af5a3>.

⁹¹ In the UK, around 88% of data breaches are due to human error rather than malicious attacks, according to a report at <https://www.verdict.co.uk/uk-data-breaches-human-error/>.

compete to maintain the ledger, it becomes extremely expensive to coordinate a 51% attack, as this would involve buying a very large number of computers. In blockchains where not much value is stored, a 51% attack could be more feasible. On the other hand, even if such an attack succeeded, participants always have the opportunity of forking and creating a new chain, where the attack is ignored.

Figure 14: Potential points of failure for permissionless ledgers and traditional databases



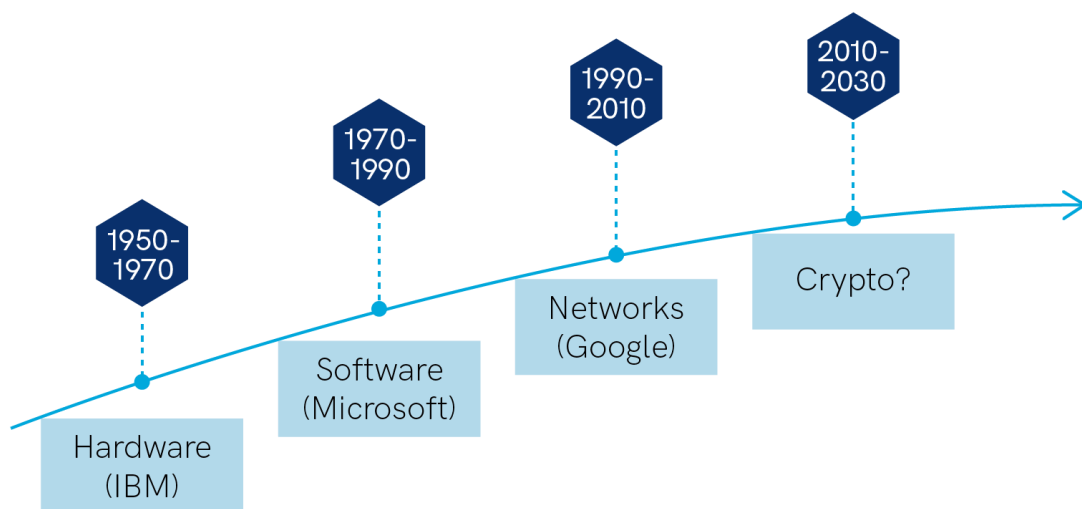
Source: Aaro Capital Research

6 The DLT Market Cycle

Distributed ledger technology is still at the early stages of development. New technologies require many years of research and development, which usually come at the expense of early investors and product creators. The following four characteristics aid new technologies in overcoming barriers to commercial adoption. First, enough use cases in the short run to motivate product developers to learn and contribute towards the technology. Second, a critical mass of people that believe in the technology, before its advantages over entrenched technologies can be realised.⁹² Third, sufficient hype and overconfident expectations, in order to attract top talent and capital that help overcome the initial high cost of development and experimentation. Finally, sufficient concentration of early profits, that incentive and allow for efficient coordination and allocation of resources, thus promoting further development and adoption. The DLT market currently satisfies these four conditions relatively well.

To put the current DLT market cycle into perspective, it is useful to compare it with previous information technology market cycles. There have been three major cycles in the past.⁹³ The first, between 1950-1970, was on hardware and led by IBM. The second, between 1970-1990, was on software and the winner was Microsoft. The networks era, between 1990-2010, is led by Google, Amazon, Facebook and Apple.

Figure 15: Past development cycles of Information Technology



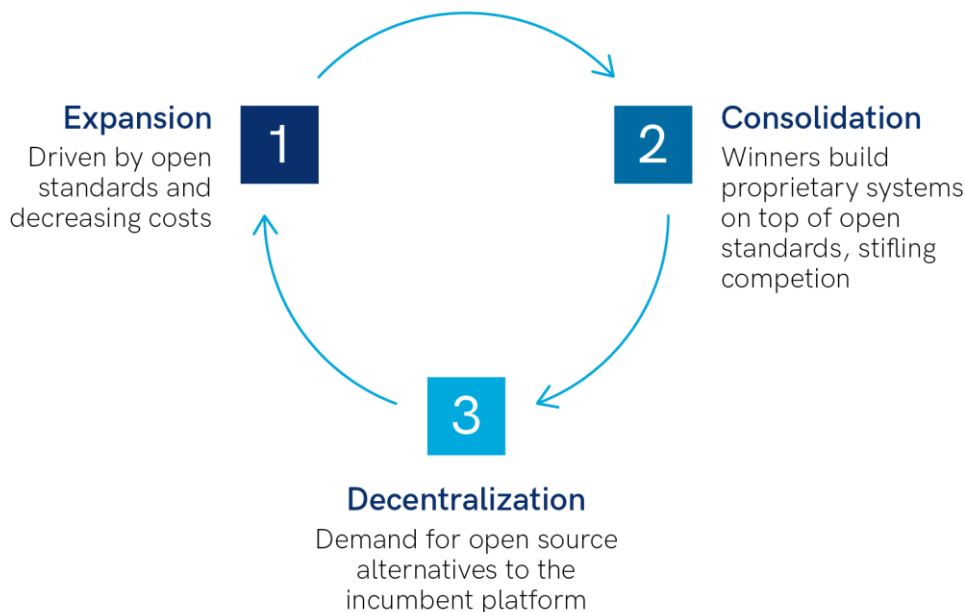
Source: Placeholder Capital

All three cycles are characterised by three phases. The first is expansion, which is driven by open standards and decreasing costs, leading to an increase in users and intense competition from start-ups. The winners lead to the consolidation phase, by building proprietary systems on top of the open standards, stifling competition. Gradually, there is demand for open source alternatives by outsiders, leading to the decentralisation phase.

⁹² For example, one reason for believing in the DLT was a deep mistrust of established institutions after the financial crisis.

⁹³ A detailed analysis can be found at <https://monegro.org/work/2018/2/20/information-technology-market-cycles-a-brief-history>.

Figure 16: Three-phase cycle of Information Technology innovations




Source: Placeholder Capital

Currently, we are at the consolidation phase of the networks era. There are very few large companies that dominate the market, each with their own proprietary systems and access to valuable user data - it is very difficult for start-ups to enter. As with previous cycles, there is an increasing demand for more consumer choice and open source alternatives. The power of the dominant network firms comes from the centralising network effects of the network layer of their platforms, with control on both the protocol and data layers.⁹⁴ Distributed ledgers are able to decentralise protocol and data ownership back to the users, while allowing network companies to still provide the same services as before. Open source and decentralised alternatives are currently being developed in the crypto ecosystem, both in terms of the infrastructure and in terms of applications. Further, there is a virtual cycle where better applications demand the development of a stronger infrastructure, which enables better applications, and so on.

To put the current technology cycle into perspective, consider the framework provided in the book “Technological Revolutions and Financial Capital”, by Carlotta Perez, which analyses five technological breakthroughs over the last 250 years, such as the Industrial Revolution and the railway boom. There are two main phases in each breakthrough. The first is the installation phase, where the infrastructure is built, and the second is the deployment phase, where the technology is broadly adopted. This is illustrated below in Figure 17.

⁹⁴ This is discussed in more detail in the paper “An Introduction to Web 3.0”.

Figure 17: Perez Technological Surge Cycle



Source: "Technological Revolutions and Financial Capital" by Carlotta Perez

Within the installation phase, there is a lot of experimentation of various technologies and the market can often misprice them, leading to bubbles that eventually burst. The 2017-2018 cryptocurrency bubble can be seen as such an episode.⁹⁵ This does not mean that it will be the last, as we have yet not entered the deployment phase, where the technology has matured, its adoption is widespread and there are a few established firms that dominate the market.

A key factor which differentiates the 2017-2018 bubble from previous ones, like the railway and the internet bubble, is that very little lasting infrastructure was created, which is nevertheless crucial for future development and widespread adoption. For example, after the railway bubble there was an overcapacity of railway tracks, which were subsequently employed by mail-to-order businesses. One of the reasons why this did not occur during the 2017-2018 bubble was the tokenisation of the asset class, allowing for capital to flow into the sector at a much earlier stage than in previous technology cycles.


⁹⁵ After a boom in 2017, the price of Bitcoin fell around 65% within one month in January 2018. By September 2018, several cryptocurrencies had decreased 80% from peak. A description can be found at https://en.wikipedia.org/wiki/2018_cryptocurrency_crash.

7 Design Issues and Solutions

7.1 The Scalability Trilemma

The design of blockchains, as platforms where smart contracts and computations can be performed, is constrained by a series of trade-offs. In a nutshell, it is very difficult (if not impossible) to achieve the following three desirable objectives simultaneously: safety, scalability and decentralization.⁹⁶

Figure 18: Impossibility Triangle for Distributed Ledger Technology



Source: Multicoin Capital

Safety refers to whether a blockchain can withstand a malicious attack, one that aims to corrupt or reverse recorded transactions. Scalability is measured by the number of transactions per unit of time that the system can perform. Decentralisation of block production (DBP) is defined as the number of independent block producers and how easy it is for a new participant to become a block producer.⁹⁷

Bitcoin sacrifices scalability in order to increase safety and decentralisation. Theoretically, it allows for maximum DBP, because anyone with a computer can be a miner. In practice though, economies of scale in mining have resulted in a few mining pools.⁹⁸ On the other hand, there have been no recorded cases of reversing a transaction that has been confirmed by at least six additional blocks. By design, Bitcoin has low scalability, as it can process very few transactions per unit of time, partly also because one block is written every 10 minutes.

⁹⁶ See <https://multicoin.capital/2018/02/23/models-scaling-trustless-computation/> for a more detailed analysis.

⁹⁷ A fourth dimension is time-to-finality, or latency, measuring how long it takes for a transaction to be considered final. For example, in Bitcoin a transaction never becomes final with certainty, but with a probability that quickly approaches 1.

⁹⁸ However, as outlined in section 4.1 **Error! Reference source not found.**, the bitcoin mining market appears to currently be competitive when measured using traditional market concentration measures.

It is important to note that, as computers become faster and more efficient, scalability increases proportionally, holding DBP and security constant. Scalability can also improve by increasing the hard-coded throughput limits of the blockchain.⁹⁹ However, if these limits increase too fast, some of the participants, who operate as nodes and maintain the ledger, may have to drop out as they are no longer able to afford investing in faster computers. If it is not easy for a market participant to maintain the ledger and verify transactions, one has to trust a third party for this function.

Delegated Proof-of-Stake protocols sacrifice the decentralisation of block production by design, in order to increase scalability and safety. One example is EOS, which has only 21 block producers, or miners, at any time, and 0.5 seconds between blocks, as compared to 10 minutes for Bitcoin. However, note that the concentration of block producers makes EOS more vulnerable to malicious attacks as compared to Bitcoin.

Other projects, such as Cosmos and Ark, sacrifice safety in order to achieve greater scalability and DBP.^{100,101} This is achieved by allowing for multiple chains that are compatible with each other. Each chain can be created easily and may support a specific application. Although each chain may be cheaper to corrupt, the value recorded on it is also lower, hence the potential gain from a 51% attack diminishes.

7.2 Layer 2 Solutions

Layer 2 solutions provide an alternative way to solve for the scalability trilemma, particularly in terms of achieving greater scalability on various dimensions. These protocol projects work by performing some computations off-chain, while still anchoring to the main blockchain to maintain security and trustlessness.¹⁰²

7.2.1 Sidechains

An important example of a Layer 2 solution is the concept of a sidechain. A sidechain is a separate blockchain that attaches to the main blockchain, as illustrated in Figure 19. The two chains communicate (sometimes in predetermined intervals), so that tokens from the mainchain are transferred to the sidechain. When the transfer is complete, computations can be performed on the sidechain, possibly using different rules and achieving different trade-offs in terms of decentralization, scalability and safety. When the computations are complete, the tokens are transferred back to the main blockchain. The mainchain only records the initial and the final states, whereas the sidechain records all intermediate states (e.g. intermediate transactions between two parties). If a dispute on the sidechain arises that cannot be resolved there, it is resolved in the mainchain by reinstating the initial state and punishing participants or redoing calculations on the mainchain (which is costly). This acts as an incentive for participants to be truthful and cooperative.

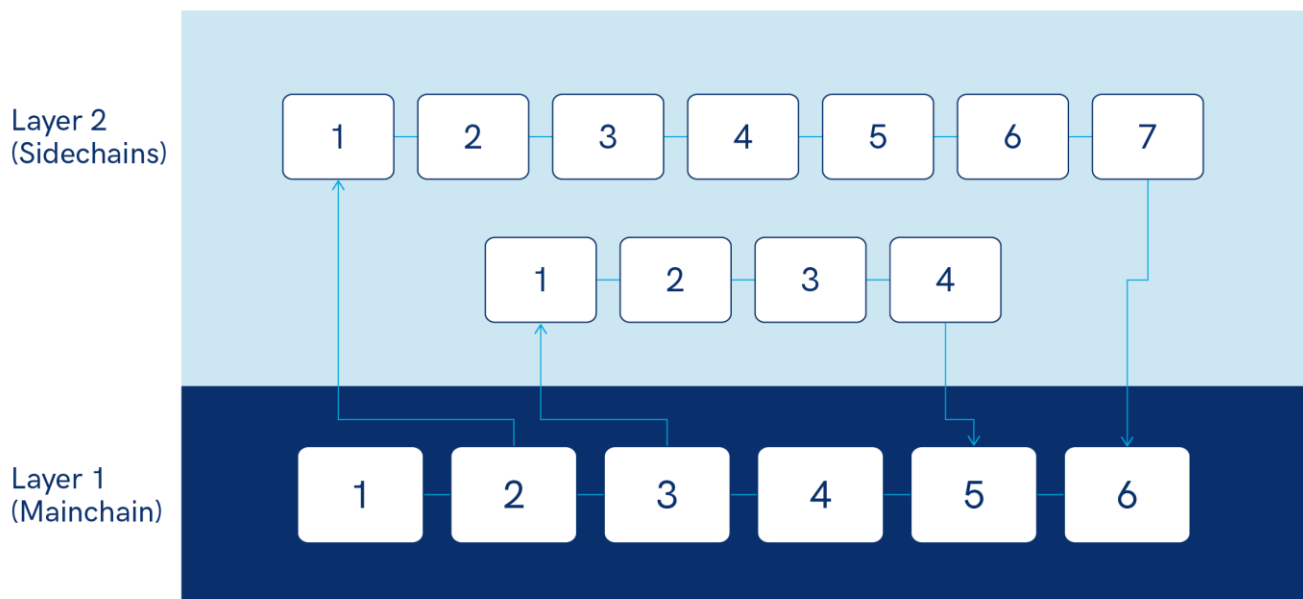
⁹⁹ For example, Bitcoin Cash was created in 2017 by imposing a hard fork on the Bitcoin blockchain. Compared to Bitcoin, it has an increased block size and can therefore average more transactions per second.

¹⁰⁰ For more information on Cosmos, see <https://cosmos.network/>.

¹⁰¹ For more information on Ark, see <https://ark.io/>.

¹⁰² An interesting analysis with several examples is provided at <https://hackernoon.com/2019-blockchain-layer-2-solution-review-d00385147396#2f47>.

Figure 19: Sidechains



Source: Aaro Capital Research

There are several sidechain projects in development. A prominent example is Rootstock, which enables smart contracts on top of the Bitcoin blockchain.¹⁰³ The Liquid Network links different cryptocurrency exchanges and traders, in order to achieve fast, private and secure Bitcoin transactions.¹⁰⁴

7.2.2 Lightning Network

Another example of a Layer 2 solution is the Lightning Network, as illustrated in Figure 20.¹⁰⁵ It is a payment network on top of the Bitcoin blockchain, enabling two users to establish a bidirectional private payment channel and then perform many transactions between them.¹⁰⁶ Transactions can settle much faster at a lower cost, since users only need to record their initial and final transactions on the blockchain. This is done by initially setting up a multi-signature wallet where each party commits some amount of BTC, and a smart contract which is essentially a payments ledger. Whenever a transaction is completed, the balance sheet of what each user owes is updated. Security can be enforced by confiscating the BTC that an uncooperative party has initially committed. When all transactions are complete, the connection terminates and the amounts on the balance sheet are recorded in the blockchain. When the network expands, users are not required to establish a direct channel with each person they want to transact with, as the Lightning Network can find an indirect path in order to establish a connection.

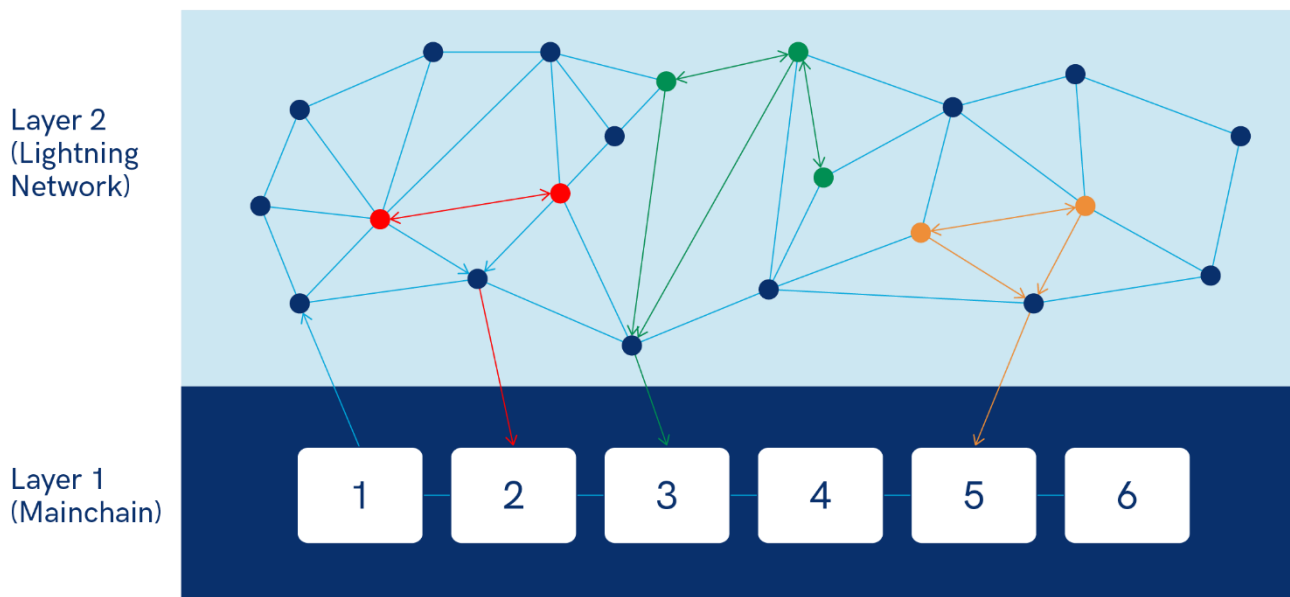
¹⁰³ For more information on Rootstock, see <https://www.rsk.co/>.

¹⁰⁴ For more information on Liquid Network, see <https://blockstream.com/liquid/>.

¹⁰⁵ The white paper can be found at <http://lightning.network/lightning-network-paper.pdf>.

¹⁰⁶ There is however a limit in the value of these transactions, which is determined by the collateral the owners of the state channel commit.

Figure 20: Lightning Network



Source: Aaro Capital Research

Payments are still executed without the involvement of trusted third parties, as they are based on a smart contract. Since fees are proportional to payments, micro-payments are possible and are settled instantly. However, if a node in the network becomes unresponsive or goes offline, the payment may be delayed or even cancelled. Moreover, large payments are not handled as effectively through the Lightning Network. The Lightning Network aims at facilitating small transactions quickly and cheaply, as compared to the more secure but slower process that involves the blockchain. A similar trade-off takes place now, where contactless payments for small transactions (below £30) are quick and cheap, but for larger transactions one has to use the more expensive and slower CHAPS payment system, which is much more secure.

7.3 Illicit Uses

One of the main criticisms against Bitcoin and cryptocurrencies in general is that they are designed to facilitate illegal behaviour: they allow for pseudonymous transactions which do not reveal the identity of transacting parties. In the early days, Bitcoin was indeed used to exchange illegal goods, for example in darknet markets such as Silk Road.

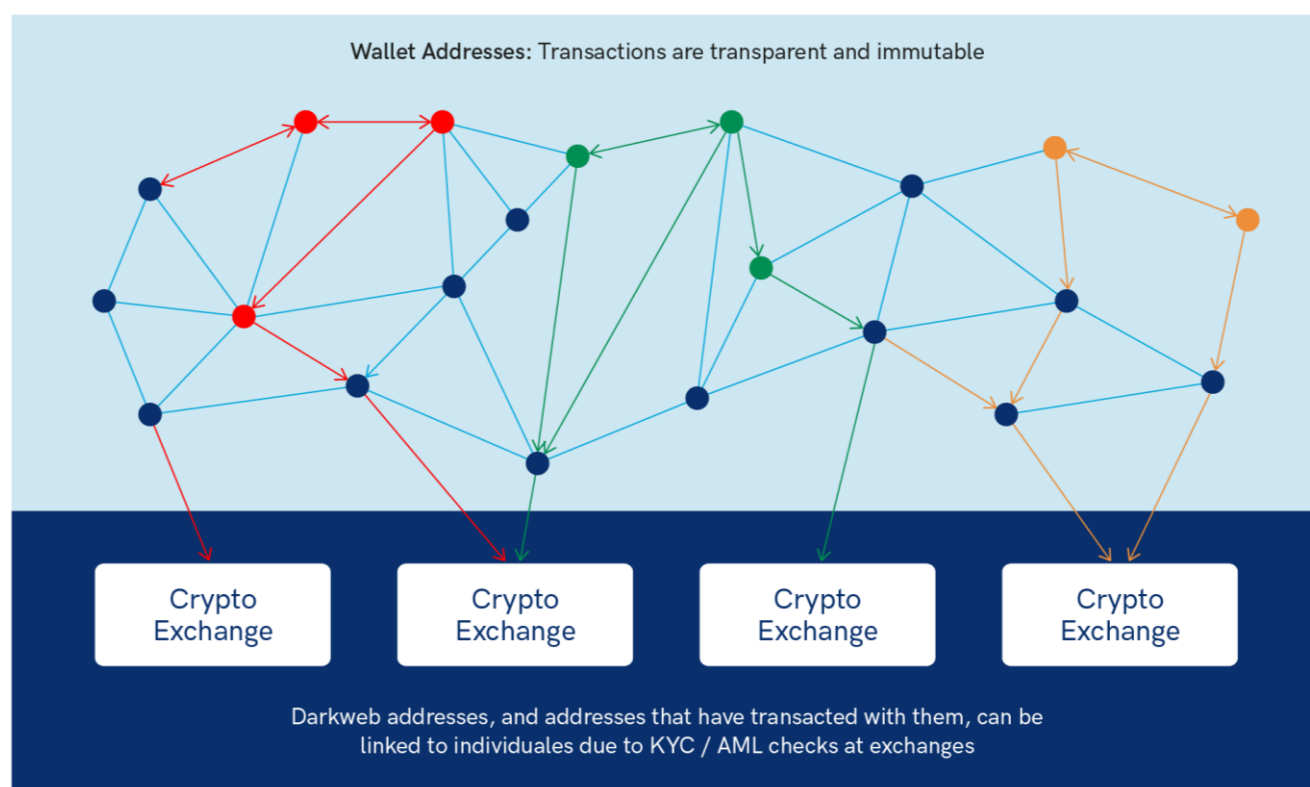
This is no longer true. A report by Chainalysis shows that the share of value in BTC sent to darknet markets has declined from 7% in 2012 to less than 1% in 2018.¹⁰⁷ There are two reasons for this. First, regulation has been updated and law enforcement authorities have started to act against these cases. For example, in June 2018 the US Department of Justice announced the arrest of 35 individuals for selling illicit goods, and confiscated nearly 2000 BTC, worth around \$20 million.¹⁰⁸ Second, the design of the blockchain, where transactions are public but

¹⁰⁷ The report can be found at <https://blog.chainalysis.com/reports/decoding-darknet-markets>.

¹⁰⁸ The press release can be found at <https://www.justice.gov/opa/pr/first-nationwide-undercover-operation-targeting-darknet-vendors-results-arrests-more-35>.

pseudonymous, helps rather than hinders authorities in their effort to prosecute illicit uses.¹⁰⁹ The public nature of transactions makes it easier for law enforcement to trace payments in BTC, as compared to transactions in any other traditional currency. For example, consider transactions between offshore accounts, which are private and require a court order in order to be revealed. Authorities need to have a reasonable suspicion to even request such a court order. Cash transactions are not only private but can also be completely anonymous - after they are completed it may be impossible to reveal the identity of the parties involved.¹¹⁰ The pseudonymity of BTC transactions is generally not a major issue for law enforcement, because once a public transaction is deemed suspicious, authorities need only to trace the conversion of BTC to a traditional currency. This is relatively easy now, as crypto exchanges are now required to follow traditional Know-Your-Customer and Anti-Money-Laundering rules to record the identities of their users. Hence, it can be relatively straightforward to uncover the identity of parties involved and subsequently prosecute.

Figure 21: Linking BTC transactions to individuals



Source: Aaro Capital Research

¹⁰⁹ There are technologies such as payment mixers with obscure the source of payments between users by mixing them with other payments.

¹¹⁰ Privacy coins, such as Zcash and Monero, operate without public addresses, thus making transactions more difficult to track. If their use proliferates and it becomes apparent that they facilitate illegal behaviour, they may be at risk of enforcement action by the authorities.

Authors:

Dr. Spyros Galanis

spyros.galanis@aaro.capital

Peter Habermacher

peter.habermacher@aaro.capital

Contact Information:

Peter Habermacher

peter.habermacher@aaro.capital

Ankush Jain

ankush.jain@aaro.capital

Sebastien Jardon

sebastien.jardon@aaro.capital

aaro.capital