Blockchain Technology:
A Guide to Analysis, Design, and Implementation

Researched and written by Alex Knight

Published by the Hexa Foundation, March 2019

HEXA
FOUNDATION

About The Hexa Foundation

The Hexa Foundation is an not-for-profit organization focused on using blockchain to create social impact. It is part of the Orbs Group of Companies and was co-founded by Netta Korin, one of the founders of Orbs, the hybrid blockchain. The Orbs Group is the largest group dealing in blockchain solutions in Israel, with over 60 employees focused on the blockchain field. The Hexa Foundation focuses on research, education and special projects which aims to use blockchain for social impact and harness the mind power of our ecosystem and network to help solve the region's and the world's most pressing humanitarian problems.

Netta comes to the Foundation following years of experience in business, government and non-profit industries. Most recently, Netta worked as a Senior Advisor in the Israeli Ministry of Defense to General Yoav (Poly) Mordechai, Head of CoGAT, and has an in-depth knowledge of the socioeconomic problems in the Gaza Strip. Prior to that position, Netta worked as a Senior Advisor to Deputy Minister Dr. Michael Oren in the Prime Minister's Office in Israel, focusing on Palestinian issues. Netta has held board positions in several non profit foundations in both Israel and the United States, and spent multiple years managing hedge funds.

For more information please contact Netta Korin (netta@hexa.org)

About the Author

Alex Knight is an MBA Candidate at the Stanford Graduate School of Business and recently ended his BTA Fellowship at New America, a technology-focused think tank based in Washington DC. Prior to matriculation and his fellowship, Alex was an investor at Tiger Management and then Blue Ridge Capital, where he focused mostly on analyzing and investing in public and private media, software, and internet companies. He graduated summa cum laude from Yale University.

HEXA
FOUNDATION

Preface: The Potential of Blockchain Technology

Since the mysterious Satoshi Nakamoto published the white paper explaining the Bitcoin blockchain in 2008, blockchains have taken the world by storm. Initially, the technology, and particularly the use of Bitcoin, was associated with the online drug trade and ransomware. However, the technology's merits led members of other sectors of the economy to further research it and consider blockchains as a means to solve protracted problems. It has since spread far and fast, with blockchain pilots of varying stages of maturity implemented across most of the major sectors in the global economy.

There are now over 1,000 different blockchain protocols and applications that are "live", and initial coin offerings (ICOs) raised close to $22 billion in 2018 alone, up from $6.6 billion in 2017 [1]. These capital inflows and traditional venture capital investments have provided the industry with the capital necessary to fund its growth. However, that same enthusiasm has generated the opportunity for disinformation to spread and for bad actors to take advantage of the uninformed.

The goal of this paper is to improve our readers' understanding of this nascent technology in the hope that doing so will aid in the development of solutions that can be used to create value in the public, private, and non-profit sectors of the global economy.

The paper is broken into six sections:

1. An introduction to blockchains, which will try to simplify the mechanisms by which blockchains function and clarify some of the industry's taxonomy
2. A review of the strengths and weaknesses of blockchains relative to their traditional database peers, with which we believe they should be compared
3. The provision of a framework for organizations to use when considering whether a problem may be suited to a blockchain solution
4. A few principles for blockchain design and deployment
5. Four case studies showing how blockchains and DLT have been deployed in different contexts
6. A summary and closing section

As this paper will show, blockchains and DLT have a lot of potential, but they are not a panacea. Contrary to what blockchain maximalists and crypto-anarchists say, though it may certainly play a more prominent role, code is unlikely to ever totally replace traditional law. Further, despite the progress being made, blockchains and other forms of DLT are still in early stages of their development. Our research suggests that rather than replacing existing systems, blockchains are likely to supplement them for the next few years. That said, the technology has opened up opportunities that were previously impossible and could conceivably create billions of dollars of value across sectors as it matures.

[1] "Coinschedule - Cryptocurrency ICO Statistics," accessed Jan 14, 2019, https://www.coinschedule.com/stats.html.

HEXA
FOUNDATION

The financial benefits of a large-scale deployment of the technology will likely take many forms, including: the disintermediation of low value-add middlemen, reduced transaction fees, improved margins, and higher liquidity in asset classes like real estate. It also has the potential to make parasitic economic activities, like fraud, usury, and exploitation, more difficult, thereby reclaiming some portion of the related ill-gotten gains.

Encouragingly, our research suggests that the value created by the technology has the potential to accrue not only to the world's largest corporations and governments, but also to the billions of people who have been largely left behind: over 2 billion people worldwide are still unbanked[2], and about 1.1 billion of them live without any form of valid identification[3]. By reducing the cost of verification and auditing and facilitating coordination, blockchains could significantly increase the participation of individuals that belong to these marginalized groups in the global economic system.

We are the first to recognize that we may well be wrong: the technology's evolution may accelerate further and begin to meet the expectations of the maximalist contingent, or it may be replaced by something new and follow in the steps of laserdiscs and bubble memory. However, both blockchains and distributed ledger technology (DLT) in general show great promise, and we hope that this paper will serve as a means to facilitate learning for technologists and luddites alike. A better understanding of the technology will increase the likelihood that blockchains or something like them will precipitate positive change by increasing accountability, reducing friction, and facilitating a more inclusive world.

[2] "2 Billion People Worldwide are Unbanked – here's how to Change This," accessed Mar 25, 2018, https://www.weforum.org/agenda/2016/05/2-billion-people-worldwide-are-unbanked-heres-how-to-change-this/.
[3] The World Bank Identification for Development Initiative, *Identification for Development (ID4D) Global Dataset.* The World Bank, 2015.

HEXA
FOUNDATION

Table of Contents

HEXA
FOUNDATION

Executive Summary

Blockchains have gone from fringe technology to global phenomenon in less than a decade. Heightened interest has been driven by the rise in the value of cryptoassets like bitcoin, the spread of claims that code will replace many of the world's institutions, and the potential the technology has to create billions of dollars of value.

Despite the widespread interest in blockchains, they can be hard to understand. While trying to avoid the industry's affinity for acronyms and technical terms, we will try to show that the brilliance of blockchain technology lies in its relative simplicity: a blockchain is a single record of transactions and/or data that is shared by multiple different parties and is very hard to change. The most important thing that makes blockchains different from traditional solutions is that they can facilitate trust and coordination without the need for a third party (like a bank). Instead of using an intermediary, other members of the blockchain network validate data according to a predetermined set of rules, which members implicitly agree to when they join the blockchain network.

None of the major component parts that make up a blockchain are new. The technology combines a digitized form of majority vote, incentive mechanisms based on centuries-old game theory, and the same cryptographic principles that underlie the vast majority of digital communications today. What is innovative is the synthesis of these elements to create a decentralized record unlike any relational or non-relational database we are aware of.

This synthesis provides blockchain technology with a number of strengths relative to its traditional peers; among them: their tamper resistance and trust facilitation, security, fault tolerance, and openness.

As a result of the differences between blockchain technology and its traditional peers, properly deployed blockchains and other types of distributed ledger technology have the potential to reduce friction, facilitate greater economic inclusivity, and improve accountability by streamlining existing processes and opening up opportunities that are impossible using traditional solutions. The opportunities to create value are especially exciting because blockchains are theoretically an open and democratic technology; as a result, groups with limited means can use them to tackle some of the world's most intractable problems and help those left behind by modern systems of finance and government reap significant benefits of the technology's growth.

However, blockchains can't solve every problem, and, like all technology, have areas of weakness and tradeoffs. Some of the technology's most evident shortcomings relative to its traditional peers are lower scalability, higher latency and storage requirements, reduced privacy, protocol inflexibility, and complex governance.

As a result of the tradeoffs required to attain the technology's strengths and of blockchain's early stage of development, some challenges can certainly be more easily and cheaply addressed using traditional databases. Further, traditional databases have benefitted from much more testing and development than their blockchain peers.

Consequently, the justification for deploying a blockchain or other type of distributed ledger rather than a traditional database must be strong from both a technological and economic perspective, and each solution should be tailored to address the specific problem it's intended to solve. It is therefore important for organizations to clearly set out their goals and to pursue a blockchain solution only once they understand the technology and have identified use cases where there exists a clear rationale for implementation.

This does not, however, mean that blockchains are destined for use only in fringe scenarios that traditional databases cannot or do not serve. In some cases, existing processes and the use of traditional technology are extraordinarily inefficient, expensive, or unsecure. In others, the use of blockchain technology will open up opportunities that were previously impossible.

In general, we recommend organizations only consider a blockchain as a solution in situations in which multiple parties that don't trust each other or can't coordinate and can't agree upon a suitable or cost-effective intermediary (or group of intermediaries) want to transact or share and reconcile data. Even when these conditions are met, every potential implementation should be framed in the context of existing prerequisites (like high quality digitized data and internet connectivity), and a case-specific analysis of risks, costs, and benefits to ensure that theory does not outpace practicality.

Given the complexity of the technology at this stage, we believe that few organizations are well suited to design, implement, and deploy a blockchain solution on their own. Thus while we encourage experimentation in certain contexts, we suggest that most organizations considering a blockchain seek the advice of multiple partners (technical and otherwise) and adapt existing platforms rather than try to create one from scratch.

Regardless of how blockchains are deployed, and despite our belief in the potential of the technology, our research suggests that some aspects of business and human interaction will forever remain beyond the reach of code. Thus, while we are of the view that blockchains or something like them will persist, we disagree with the maximalist claim that code will ever totally replace law. Admittedly, humanity is prone to irrationality and error, and we agree that blockchains can help reduce their prevalence; however, our foibles are part of what make us human, and thus we hope that some of them are here to stay.

HEXA
FOUNDATION

A checklist for deploying blockchain technology

- Define the problem and the goals you're looking to achieve: Blockchain isn't a panacea. As such, it is important to understand whether the technology can help achieve an organization's intended goal, and to put checkpoints in place to ensure that it is performing as intended.

- Ensure the presence of prerequisites: Aside from ensuring the technology is the right fit for the problem, organizations should also confirm that the data being input into the system is high quality and accurate, and that those intended to benefit from the system have internet access of sufficient quality for the applications that are being considered.

- Consult with potential stakeholders: talking to potential network participants and other stakeholders early in the project's life and throughout its design and implementation can help address pain points before they become impasses, leading to more efficient development and, often, improved solutions.

- Design blockchains like hardware, not software: blockchains are hard to change after they are deployed (like a physical computer as opposed to a software program, which can be more easily modified). Consequently it's important to think ahead during the design process and test extensively before deployment.

- Consult with experts: Given the nascent stages of blockchain's development, engaging technical partners and legal counsel at the outset of a project's design can help prevent technical and regulatory issues later in development.

- Maximize interoperability and avoid vendor lock-in to the degree possible: Prepare for the ecosystem's evolution by maintaining flexibility by maximizing interoperability with other systems (both blockchain-related and otherwise) and avoiding vendor lock-in where possible.

- Start small and run redundantly: Like all new technology, blockchains can experience "growing pains". These will take time to resolve. As such, we recommend starting with a pilot and, in cases where a legacy solution exists, running the blockchain redundantly with existing systems. Do not "rip and replace".

- Ask Questions: Blockchain is complex, and it's unlikely any individual knows everything necessary to successfully implement the technology. Questions can clarify problems and lead to innovation. If organizations implementing blockchain cooperate with each other by seeking input and solidifying the industry vernacular, it will benefit the ecosystem as a whole.

HEXA
FOUNDATION

Do you need a **blockchain?**

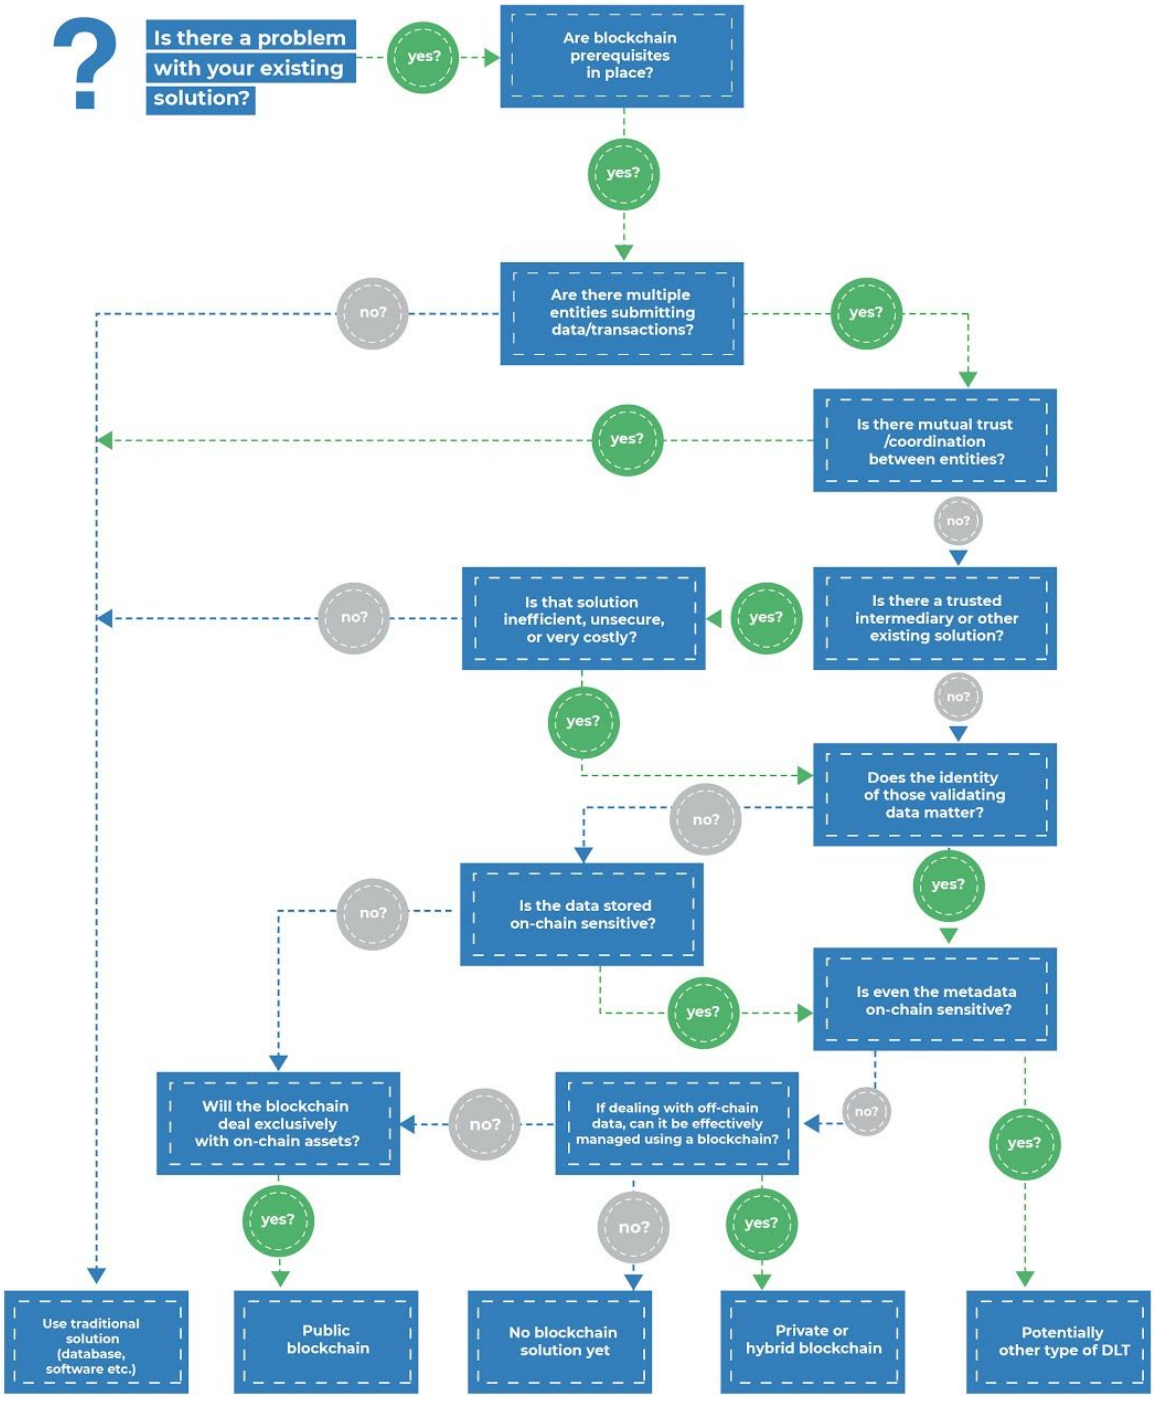

HEXA
FOUNDATION

Part I: An Introduction to Blockchain Technology

The rapid pace of the technology's evolution and technical jargon developed by the "crypto" community can make it easy to forget that the brilliance of blockchains lies in their relative simplicity. This chapter seeks to explain the mechanisms by which blockchains work in order to facilitate an understanding of concepts in later parts of this paper. It will not dig into the technical details of the algorithms and cryptography used: they are beyond the scope of this paper and are well covered by other authors.

In one sentence: a blockchain is a record of transactions and/or data that is shared by multiple different parties and is very hard to change. More technically, it's an append-only distributed database that is transparent and decentralized. Though the best-known blockchains are those related to crypto-currencies, the applications of a blockchain need not be currency-related; as we will show later in this paper, many are not. What makes a blockchain special is that the computers (known as "nodes") that constitute the backbone of the blockchain network each verify that a transaction adheres to certain rules, thereby removing the need for a central authority or trusted intermediary. Simply put, a blockchain facilitates trust and coordination between strangers. Further, because of the technology's cryptographic properties, once the network validates a transaction, that data is very hard to tamper with. The result is a blockchain's "prefix": a single list of transactions/data about which all nodes in the network come to agree.

Blockchain is a relatively new technology, but many of the concepts behind it are not. What makes blockchains innovative is their combination of these proven components in a single technology to create an accountable, secure, and efficient system.

As we will show in this chapter, there are different kinds of blockchains, but they all generally have three sets of goals in common:

A. Creating a secure network that is capable of executing, verifying, and reaching consensus on the state of shared data without the need for a trusted third-party intermediary
B. Ensuring that data can't be changed after it has been accepted by the network (known as tamper-resistance, sometimes mistakenly called immutability)
C. Verifying that the parties sending data or executing transactions have the means and authority to do so (authentication)

The mechanisms by which these goals are typically accomplished are a combination of cleverly coded protocols and cryptography. For the purposes of this paper, they can be grouped into three categories:

1. Distributed consensus in a peer-to-peer network, which allows verification of transactions and data without the need for a trusted third party
2. Incentives and/or punishments, also known as a "consensus mechanism", which are used to encourage network participants to only verify valid transactions and data
3. Public key cryptography, which facilitates authentication

The implicit result of these goals and mechanisms is that properly deployed blockchains (and some other types of distributed ledger technology[4]) have the potential to reduce friction, facilitate greater inclusivity, and improve accountability, thereby creating significant value for public, private, and social impact organizations. In some cases, we believe that they will facilitate immense improvement in the quality of life of many underserved populations and may be able to restore some of the trust in the very institutions that some authors say blockchains will completely displace.

Distributed Consensus in a Peer-to-Peer Network

Distributed consensus in a peer-to-peer network is the first critical mechanism used in all blockchains. In most transactions today, a third party like a bank, broker, or government/legal system is trusted to ensure that a transaction is authenticated and tamper resistant. Rather than rely on a third party, blockchains use a form of majority vote among a network of peers to ensure the validity of data. That majority vote is called distributed consensus.

Though the exact mechanism by which distributed consensus is achieved can differ by network, it generally follows a similar series of steps. Each node in the network checks the validity of the data it receives from other nodes, primarily by authenticating the participants involved (and verifying their ability/authority to transact or transmit the information), and by checking that there has been no attempt to break any of the network's rules or modify past data/transactions. If the transaction meets those criteria, the node in question then broadcasts it to its peers (which repeat the process) and includes the transaction in the current group of transactions it has received. This grouping is known as a "block". Each block is cryptographically linked to the blocks before it: hence the term "block*chain*".

The data stored in each block on the chain can take multiple forms and extends beyond information regarding transactions. It can be the original data (whether in encrypted or unencrypted form) or a unique digital fingerprint, that is derived from the data (using a cryptographic algorithm). When only the fingerprint is stored on the blockchain, the data itself remains in a traditional database.

[4] Please see glossary: blockchains are the best-known type of distributed ledger technology, but not the only one.

Regardless of where the original data is stored, any small change to it (even adding a comma) will yield a completely different digital fingerprint and will also change the block that contains that fingerprint. Moreover, because the blocks are linked to each-other, any change to one block also changes all the blocks that come after it, raising red flags throughout the network.

Public and Private Blockchains

Before continuing in our discussion of the critical aspects of blockchains, it is important to understand that blockchains can differ significantly along a number of lines, especially their degree of decentralization and privacy.

The best-known blockchains are public (sometimes also called "permissionless") blockchains, such as Bitcoin or Ethereum. The contents of a public chain are visible to anyone, and no permission is required to join a public chain network or to read, write, and validate the data on it. The only hurdles to participation in a public chain are an Internet connection and a computer with enough space to store relevant portions of the blockchain.

Private (sometimes also called "permissioned") blockchains are somewhat more complex and can be customized to be more (but not completely) centralized and private. There is a debate in the blockchain community as to whether private blockchains should be considered blockchains at all, the details of which are technical. What is important is that they differ in four main ways from their public peers:

Membership in the network as well as permission to read, write, and validate data is provisioned by an entity, or group of entities, already in the network; it is for this reason that private chains aren't strictly decentralized
All members of a private chain are typically identified in some way: permissioning entities need to know who they are dealing with before deciding whether to admit the entity into the network and what that entity is permitted to do after being allowed to join
Because nodes are identified, the types of consensus mechanisms (incentives and punishments explained below) in private chains tend to vary more than those in public chains
Lastly, private chains can be designed to allow anyone to read the contents of the chain or to limit legibility to private network members, allowing for greater discretion

Some chains are a mix of public and private protocols. These are known as "hybrid" chains and vary in their structure. The rules that govern hybrid chains are typically more restrictive than those that govern a public chain; however, the result of transactions that occur on the hybrid chain are always recorded on a public ledger in some way. This can be either because a private chain is directly connected to a public one (known as a "side chain") or because a fingerprint of the data on a private chain is periodically posted to a public blockchain (known as "anchoring"). As a result of either design, hybrid chains benefit, to some degree, from the public chain's technically greater security.

Consensus Mechanisms: Incentives and Punishments

Distributed consensus is reinforced by incentives and punishments that promote good behavior and constitute the second critical mechanism in a blockchain. They use game theory and codified rules to lead the network to a single, accurate, version of the "truth".

In public blockchains, these incentives and punishments are typically financial: nodes are incentivized to validate transactions due to a type of bounty awarded to the node that "mines" (finds and publishes) the next block. A successful miner earns a "block reward" (a predetermined quantity of a cryptoasset, like bitcoin, which the miner receives upon submitting a block made up of valid transactions), and/or fees paid by those transacting in exchange for the inclusion of their data in that block of information.

The only rewards recognized by a blockchain network are those associated with the creation of blocks included in the longest version of the chain, known as the "prefix". Consequently, if a proposed block is unlikely to be part of the prefix for any reason (especially if it contains an invalid data), miners are incented not to link subsequent blocks to it, leaving it "orphaned". The result is that, while it may take some time, nodes converge on a single version of the data that they agree is compliant with the network's rules.

There are several different consensus mechanisms to determine which miner will be chosen to propose the next block in a public chain. The most common of these are known as "proof of work" and "proof of stake." While we will not cover their technical aspects, what is important is that these consensus mechanisms require miners to commit significant resources to the chain by investing in either computing power (in the case of proof of work), or a large quantity of the cryptoasset they intend to mine (in the case of proof of stake). The value of those resources depends in part on continued collective trust in the blockchain: if other network members feel that those mining blocks are violating the rules, the other users will leave the network and the value of all miners' investments in that network (hardware and/or cryptoasset holdings) will decline. As a result, miners are further incentivized to only include transactions that comply with a network's protocols in their blocks: if they don't they'll lose money.

Equally important is that consensus mechanisms introduce a degree of randomness as to who will propose the next block: the probability of a miner being chosen to propose a block is tied to the relative amount of network computing power or cryptoasset (token) holdings that the miner controls. Over time, the share of blocks a particular miner in a public chain is chosen to propose should equal that miner's share of that resource.

Consequently, if an attacker wants to ensure their invalid block isn't rejected by the network, or to change data in a previous block, they would need to make a huge investment of resources in order to have more computing power or cryptoassets than the rest of the network combined [5]. So long as good actors in a network control the majority of computing power/cryptoassets, the chances that an attacker can permanently pass off an invalid transaction, change a past

[5] There are some theoretical examples in which a bad actor could dominate a chain with slightly less than 51% of the computing power or asset supply

transaction, or censor a piece of data, become increasingly negligible as more blocks are appended to the chain[6].

Because there is more trust inherent in a private chain than in a public one (nodes in a private chain are typically identified before being allowed to validate transactions), the consensus mechanisms in private chains vary broadly and can depend on off-chain interactions and presumed investment in the success of the system. For example, in a supply-chain-focused blockchain, a validator may be incented to be honest due to the threat of retaliation from other members of the network with whom they interact in the real world. In the case of a consortium, members may presume their peers benefit from the success of the protocol and want to preserve their reputation. In any case, private blockchains are also tamper resistant to varying degrees and eventually arrive at an agreement on the state of shared data.

Public Key Cryptography

As noted above, part of the validation process requires establishing that the actors transacting from specific accounts are authorized to do so. This authentication is accomplished using public key cryptography, which is the third critical mechanism in blockchains.

Public keys are the equivalent of username. Like usernames, they can be changed frequently and are often pseudonymous: just because the network knows a user name/public key doesn't mean it knows the identity of its owner. Each public address is linked to a set of permissions (if it is on a private chain) and to a balance or "state," which is typically visible to all network participants and denotes the resources that the owner of the address can spend. Individual entities in public chains can have multiple public keys, create new ones, and move resources from one address to another with virtual impunity.

Each public key is linked to a private key, which is equivalent to a password: only the owner of the public address should have it. The private key generates a digital signature that is cryptographically linked to its corresponding public address and to the data that the signature is appended to. Therefore, the presence of that signature announces to other network participants (who are validating the data or transaction) that the address' owner has authorized the contents without revealing the information or the private key itself.

Thus, public key cryptography combines with distributed consensus and consensus mechanisms to produce a dynamic that allows all actors in a properly functioning blockchain to safely transact without having to trust their counterparties or rely on an intermediary. Each block (made up of data and/or transactions) produces a shared record and is linked to one or more addresses and to the previous blocks in the chain. Because of these links and the chain's consensus mechanisms, the entire network converges on a single version of the blockchain that can be audited by all members of the network.

[6] Satoshi Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System" Bitcoin.org, 2008). 6.

HEXA
FOUNDATION

The Implications of Blockchain's Growth: Reduced Friction, Increased Inclusivity, and Improved Accountability

Regardless of their form and function, the widespread implementation of blockchain technology that we have reviewed in our research would have three implicit results with meaningful implications: reduced friction, increased inclusivity, and improved accountability.

Friction is the "costs, both implied and direct, associated with a transaction"[7]. These costs include everything ranging from transaction fees and investment in research to the expenditure of time and effort. Blockchains have the potential to meaningfully reduce friction in multiple ways, including lowering fees and processing times, automating execution, and diminishing human error. While blockchains are unlikely to outright replace legacy systems in every case, they can create considerable value by supplementing existing processes and creating new opportunities.

The lower cost of verification and auditing associated with blockchain technology mean that it can facilitate the coordination between entities that would otherwise be impracticable or economically infeasible. In certain configurations, blockchains could allow institutions to process data and transactions for customers that it is currently uneconomical for them to reach and serve. As a result, the technology offers a means for organizations across sectors to better facilitate increased financial inclusivity and help many of the world's underserved.
Lastly, blockchains also have the potential to meaningfully improve accountability and transparency. If properly implemented, they could provide a means to track assets and resources consistently, leading not only to reduced corruption, loss, and waste, but also to better decision-making based on objective evidence that is available to all stakeholders for review and analysis. As such, while they may displace some institutions through decentralization, blockchains can also provide a means for other organizations to recover trust by implementing this transparent new technology.

[7] "Definition of "Friction Costs" - NASDAQ Financial Glossary," accessed Apr 17, 2018, https:/www.nasdaq.com/investing/glossary/f/friction-costs.

HEXA
FOUNDATION

Part II: What makes Blockchains Different?

Despite their technical promise and meteoric rise in the public consciousness, blockchains are early in their evolution; as such, some problems can be more easily and cheaply addressed using traditional databases. There are, however, areas in which a blockchain can offer significant advantages over its traditional peers and where the novel technology will open up entirely new opportunities. Understanding what makes a blockchain different and why that matters is of critical importance when deciding whether or not to deploy the frontier technology. The goal of this chapter is therefore to address some of the technology's major areas of strength and weakness relative to traditional solutions and their implications.

Strengths:

- Tamper resistance and trust-facilitation
- Censorship resistance
- Openness
- Security
- Fault Tolerance

The strengths that blockchains have relative to traditional databases come at a price: the technology is also underperforms traditional databases in certain areas. It is important to understand that trade-offs like these are typical of any form of technology, not just blockchain. Not all databases are a match for every goal, and diverse use cases will require design choices that will lead to variation in how a solution performs.

Weaknesses:

- Scalability and latency
- Storage requirements
- Privacy
- Inflexibility and complex governance
- Ensuring data quality
- Unpredictability of consensus mechanisms, and protocol sustainability

Further, though we believe that the characterization of the technology's strengths and weaknesses herein is accurate, our evaluation is made on the basis of blockchains as they exist today, which is to say: very early in their evolution. Given the immaturity of the technology, it is likely that as blockchains and traditional databases evolve, some of their relative strengths and weaknesses will be addressed or eliminated.

Strength: Tamper Resistance and Trust-Facilitation

As described in Part I, what makes blockchains different from traditional databases is they have multiple entities that can read/write to them and cryptographic consensus mechanisms that all nodes use to decide whether to consider data valid. This allows blockchains to facilitate trust without relying on a centralized intermediary.

These particularities produce two big benefits relative to a traditional database: tamper resistance (sometimes called immutability) and the ability to facilitate "trust across a boundary[8]", be it organizational, geographic, or otherwise. This quality is sometimes also referred to as "trustlessness".

While tamper resistance is a feature common to all blockchain types, the degree to which a blockchain is trustless depends on the type of blockchain and on the rules of a particular protocol. Public blockchains, like Bitcoin, are theoretically more trustless and decentralized. Trust in the system arises from the financial incentives present for miners to only include valid transactions in their blocks. As long as "good" actors control the majority of computing or cryptoasset resources in a network, and protocols function as expected, public chains are completely trustless.

The degree of trustlessness in private blockchains varies from case to case and depends mostly on the consensus mechanism used. No private blockchain is completely devoid of trust in that there is an entity or group that selects who is allowed to be part of the network as well as what new network members are permitted to do (read, write, validate etc.) In most cases, the differentiated authority of permissioning entities ends there: once members are admitted, predefined protocols apply equally to all nodes in the network, whether they are permission-providers or not.

Strength: Censorship-Resistance

Another advantage of decentralization is that no single entity can prevent a piece of data from being posted to the blockchain. This is known as censorship-resistance. Like tamper-resistance, the degree of censorship-resistance depends on the type of protocol and the rules that govern it. If they have achieved scale, public blockchains are highly censorship-resistant, while private chains can vary depending on their rules.

Strength: Openness

Decentralization also contributes to the fact that blockchain networks tend to be more open than traditional databases are, meaning that they are easier to participate in. In public chains, anyone can participate in the protocol without the need for a central entity to coordinate or approve activity: all that is required is an Internet connection and a computer (a smartphone can serve in certain cases). As we will elaborate on later in the paper, this characteristic pairs with the trust-facilitating nature of blockchains to enable coordination among entities that would otherwise be impracticable.

Depending on the rules of a particular protocol, private chains can also be more open than traditional databases are: network membership can often be approved by any of a number of entities that are already part of the network. That number is usually greater than the limited number of permissioning entities that allow participation in the case of a traditional database.
We note that while a blockchain may be more open than the typical database, the on-ramps to the chain (such as specialized hardware, wallets, and exchanges) are often not.

[8] Bob Visnov, Conversation with Bob Visnov, Apr 5, 2018.

Consequently, participation in blockchain networks can be limited either through regulation of these on-ramps or through the control of Internet access (among other methods). Further, while participation in the protocol may have a relatively low bar, mining (and therefore validation) on public chains in particular, often requires the investment of large amounts of capital and considerable technical knowledge. This can limit the number of entities that have meaningful input into a blockchain's governance.

Strength: Security

There exists no single point of vulnerability in a blockchain that an attacker can target in order to "hack" the network. As a result, blockchain networks are typically more secure than traditional databases are.

The more decentralized and large a blockchain network is, the more secure it theoretically is. Consider a properly functioning public chain, where an attacker would have to control over 50% of the computing power / network cryptoasset holdings before they are able to change any aspects of the protocol or modify past transactions. This is meaningfully more difficult than compromising a single user's computer (as when attacking a traditional database).

That said, not all elements of the blockchain ecosystem are invulnerable: despite the security of the core technology, blockchain users are currently no less susceptible to the loss or theft of private keys (which are effectively passwords) than are users of traditional databases. Further, the larger number of database instances (data is stored by each node) on a blockchain actually provides more potential targets if the attacker's only goal is just to view the information stored on the database rather than compromise the network as a whole. Consequently, if network members want information to remain highly confidential, special measures should be taken and a blockchain may not be the best solution.

Furthermore, many of the organizations that serve public blockchain users (like wallets or exchanges), lack the decentralized characteristics and consensus mechanisms that make the core blockchain secure. As evidenced by episodes like the Mt.Gox and DAO hacks in the Bitcoin and Ethereum ecosystems respectively, while the blockchain may be highly secure, the services and applications that sit on the platform may not be. Lastly, it is important to recognize that because the value of a blockchain is in part based on the size of its network, an attack that leads to a loss of trust in a protocol, even if it has little fundamental impact on its operation or security, can be devastating.

Strength: Fault Tolerance

A less intuitive benefit of blockchains is that they are generally highly "fault tolerant." Fault tolerance refers to the ability of a database to function despite the failure of multiple nodes in its network. Depending on the structure and size of the network, if some of the nodes in a traditional database malfunction, the entire network can crash. Conversely, because each full node on a blockchain network stores a copy of the whole database, blockchains are theoretically much less likely to "go down". Blockchains are therefore one means of ensuring the reliability of critical applications and the survival of data in the event of a catastrophe.

We note that while blockchains theoretically have higher fault tolerance than their traditional peers, that assertion relies on a number of assumptions about both the blockchain in question and the traditional alternative. In order to be highly fault tolerant, a blockchain must have a large number of full nodes on the network; because there is no single entity in control of a blockchain, this is not guaranteed. A blockchain can't order its network participants to add another node; the protocol must incent them to do so in some way. Further, traditional databases can also be structured to have very high "redundancy" by having copies of each piece of data on many separate nodes (as is the case on many modern cloud storage systems), thereby effectively creating the same benefit.

Weakness: Scalability and Latency

While blockchains provide a number of very attractive qualities, those features result in a number of drawbacks that organizations should consider. The most widely recognized of these issues are that blockchains are less scalable and slower (have higher latency), than traditional databases are. These problems arise as a result of the decentralized and secure nature of blockchain technology and can be mitigated, but not eliminated.

Scalability "connotes the ability of a system… to process growing volumes of [information]"[9]. In order to allow communication across the decentralized peer-to-peer network in a blockchain, multiple nodes typically validate each piece of information. As such, the chain is only as scalable as its full nodes are. When a traditional network is limited by the computing power of its nodes, IT teams can increase the amount of power available by installing extra servers or modifying network software. However, as we noted, no centralized entity controls a blockchain. Thus, the amount of computing power dedicated to a blockchain network is generally determined by the incentives or punishments associated with a particular protocol rather than what the network actually needs in order to cope with growing data volumes.

The scalability problem is exacerbated by the protocols of many blockchains, which impose limits on transaction volumes and block sizes to enhance security. The Bitcoin protocol, for example, is designed so that blocks are formed about every 10 minutes and can only be of a certain size (1MB), which is roughly equivalent to 500 typed pages. This is true regardless of the volume of information being processed and further limits the maximum throughput of the network.

These protocol-based limits also produce much higher latency (how slow the network is): transactions are processed and committed to memory not as quickly as is possible, but as the protocol dictates. This leads blockchain networks to cope with information more slowly than a traditional database could.

As more nodes are added to a blockchain network, latency issues are exacerbated: there are physical limits to how quickly information can travel across a network, and introducing more nodes increases the time it takes for data to propagate across database instances. While the latency differential between blockchains and their legacy peers varies significantly depending on the particular chain, consider the following: Ethereum's network currently processes around

[9] André B. Bondi, *Characteristics of Scalability and their Impact on Performance* (New Jersey: AT&T Labs,[2000]).

16 transactions per second (TPS)[10], Visa's network is capable of over 24,000 TPS[11], and in 2010, Facebook's database could process up to about 13 million queries per second[12].

There are certain compromises (usually giving up some degree of security in exchange for greater scalability and speed) that can be made to mitigate these shortcomings, and improving them is an area of research that has attracted many of the brightest minds in the blockchain world. As a result, we expect that blockchains will become more scalable and fast over time, though they are unlikely to ever match a well-designed traditional database's capabilities.

Weakness: Storage

While blockchains' decentralization makes them very robust databases, it also makes them a highly inefficient means of storing large amounts of information. Each full node on a blockchain network stores a full copy of the database. As a result, the amount of storage required by the network increases exponentially with each piece of data added to the blockchain: on January 14th 2019, there were about 9,860 full nodes on the Ethereum blockchain[13]; thus for every additional 1MB stored on Ethereum, the network as a whole required close to 9.7GB of additional storage space.

In the technology's defense, nodes don't necessarily need to have every piece of data to validate transactions: blockchains can serve in the capacity of an audit trail for data that is stored in a traditional database off-chain, or a blockchain can implement some form of database "sharding", whereby some nodes only store data that is relevant to them. That said, the decisions involve tradeoffs that include somewhat reducing the fault tolerance and security of the network.

Weakness: Privacy

One of the frequently cited characteristics of blockchains is that they are more transparent than their traditional peers: with the exception of some public blockchains specifically designed to provide anonymity, anyone with access to the network and appropriate permissions can see each block on a blockchain and the transactions that make it up. While it is true that many blockchains are pseudonymous and that data is usually encrypted, there are a number of methods available that allow a dedicated observer to link transactions on many blockchains to real world identities. Further, in some scenarios, metadata (data that describes data, like when a transaction happened) may, itself, communicate sensitive information. Organizations can increase the difficulty of transaction analysis by using a private chain, on which the ability to read the contents of the chain can be restricted.

[10] "Ethereum Charts and Statistics,"accessed Feb. 26, 2018, https://etherscan.io/charts.
[11] "Small Business Retail - Visa," accessed Feb 28., 2018, http://usa.visa.com/content/VISA/usa/englishlanguagemaster/en_US/home/run-your-business/small-business-tools/retail.html>.
[12] *MySQL Tech Talk - 11.2.10,* Facebook Live, directed by Mark Callaghan (Menlo Park, CA: Facebook, 2010)
[13] "Ethernodes.Org - the Ethereum Node Explorer," accessed Jan 14, 2019, https://www.ethernodes.org/network/1.

However, those trying to join a private network must provide permissioning entities (usually one or more current members of the network) with aspects of their identity in order to be added to the network in the first place. This, in turn, may produce its own privacy-related concerns and reduces the degree of decentralization of the chain as a whole. Alternatively, users may choose a privacy-focused protocol, thereby opting for greater discretion in exchange for higher latency and lower scalability.

Given that the methods used to encrypt information on a blockchain are typically the same as are used in traditional systems, many in the blockchain community argue that greater transparency is a small price to pay for improved security. Very few actions we take on the internet (or in life) are truly private as it is, and the organizations we trust with our information often do a poor job of keeping it secure. Further, many companies in the blockchain community are trying to utilize the technology to improve individual privacy by providing a means to attain data sovereignty. However, the fact is that one of the features of most blockchains is transparency, and if data or its metadata are particularly sensitive, blockchains are not as private as a secure traditional database is.

Weakness: Inflexibility and Complex Governance

The decentralized structure and governance of blockchain protocols makes them harder to modify than traditional databases. Changing the rules of a blockchain typically requires the agreement of the majority of the other stakeholders in a network (of which there may be many), and some aspects of protocols cannot be changed at all.

Depending on one's perspective, this inflexibility can be a strength or a weakness. Many organizations adopt blockchains precisely because they are hard to change; however, sometimes change is necessary. For example, if a social impact mission changes, or a bug in a blockchain's code is discovered, then it may be desirable to change the underlying protocol. Further, blockchain technology itself is evolving extremely quickly, and the inflexibility of protocols today may make interoperability with new protocols challenging.

It is true that traditional databases can also be hard to modify from a technical perspective (though governance is much easier), and that members of public blockchains can implement significant changes to their protocols through a "fork". Forks are a name for protocol splits that allow network members to choose to build on whatever chain has the set of rules they prefer. However, these forks will likely diminish the value of the chain as per Metcalfe's Law[14]: the value of a network increases (and decreases) exponentially with its number of nodes[15].

[14] Christopher S. Yoo, *Moore's Law, Metcalfe's Law, and the Theory of Optimal Interoperability* Penn Law: Legal Scholarship Repository,[2015]).
[15] We make this argument with the generalized blockchain concept in mind. Some would point out that in the case of cryptocurrencies, it may be users, not nodes, that help define the value of the network.

HEXA
FOUNDATION

Weakness: Ensuring Data Quality

High quality data is key to the trust in a blockchain's consensus, and blockchains have no way of natively determining the validity of data generated outside of the network. All inputs in on-chain transactions can be audited and verified by full nodes according to the chain's consensus mechanism. However, the blockchain community has not yet developed an effective means to reliably link on-chain data to off-chain reality (digital or physical). There are promising developments in this area (such as oracles and token curated registries, to name two). However, blockchain network participants will likely have to trust traditional mechanisms to some degree to derive off-chain data (in the same way traditional database users do), and at times resort to traditional means of recourse (like filing a lawsuit) if/when that data proves inaccurate.

Weakness: Unpredictability of Consensus Mechanisms and Protocol Sustainability

Public blockchains like Bitcoin and Ethereum are (at least theoretically), more decentralized and secure than their private peers. As a result of the fact that there is no need to be "admitted" to the network, these chains are also more open and censorship-resistant.

One of the tradeoffs is that public chains rely on consensus mechanisms based on financial rewards provisioned to the miners who verify transactions. This in turn relies on the perceived value of a chain's cryptocurrency/tokens (like bitcoin in the Bitcoin blockchain) and on the rules that determine the nature and magnitude of those rewards. There is no guarantee that these rewards and the game theoretical mechanisms that support them will continue to drive optimal behavior by miners. If they don't, the consensus mechanism that supports the chain would collapse, leading to a mass exodus of users from the network as trust in the blockchain in question evaporates.

From a different sustainability perspective, one of the most commonly used consensus mechanisms, known as Proof of Work, requires the use of massive amounts of computing power. It is therefore highly energy intensive. As a result, unflagging growth of Proof of Work based protocols is likely to have significant negative implications for global energy consumption.

Unlike public blockchains, private chains generally don't rely on the financial incentives that drive miners in a public chain. The incentives for acting honestly in private chains vary and are typically not underwritten with the same degree of game theoretical care that those on many public chains are. The conclusion is that the longevity of their incentive mechanisms and, by relation, of early-generation private chains, is no more guaranteed than that of their public brethren.

Organizations considering blockchain technology should therefore think about the immaturity of incentives, platforms, and the ecosystem as a whole. The Internet has changed considerably since 1990; while we don't believe that blockchain and the Internet are completely analogous, there are still many more unknown than known elements in the blockchain ecosystem, as there were in the early days of the internet. As such, many of the companies that were leaders in the early days of the Web's evolution no longer exist. The same may be true of early-generation blockchain protocols. We are so early in the evolution of the technology and specific protocols, that predictions regarding which platforms will persist may prove inaccurate.

Conclusion

While the aforementioned analogy to early Internet development is intended partly to serve as a caution, it should also be viewed as evidence of the possibilities of a nascent technology.

The concept of the Arpanet was developed in 1965 and it wasn't until 1990 that the "world wide web" was introduced to the public. Yahoo! launched its website in 1994, and it took another four years for Google to be founded[16]. Further, the Internet had a lot of issues early in its infancy, and still has many today: it was slow, it wasn't particularly scalable, and it's still plagued by security concerns (some of which blockchain may be able to resolve). Even after nearly three decades of evolution, the Internet and its associated technologies have failed to solve every problem; yet it has unquestionably revolutionized the way the world works and helped billions of people.

Our belief is that the blockchain ecosystem will take time to develop in the same way that the Internet has. Consequently, though the tradeoffs associated with blockchains may seem significant today, we believe that Distributed Ledger Technology (DLT) is here to stay, that there are a number of applications for which it makes sense today, and that that number will grow as the technology matures. The community is working to address many of the weaknesses denoted above, and while it may not eliminate them all, there is little doubt that blockchains and their peers will evolve meaningfully over the next five to ten years.

[16] "Internet History Timeline: ARPANET to the World Wide Web," accessed Mar 14, 2018, https://www.livescience.com/20727-internet-history.html.

HEXA
FOUNDATION

Part III: Do you need a Blockchain?

Traditional databases are much further along in their development than their distributed ledger peers. As a result, they have undergone more testing, have developed more refined programming languages, and benefit from a larger community/ecosystem. Consequently, the justification for deploying a blockchain or other type of distributed ledger rather than a traditional database must be strong from both a technological and economic perspective. Further, blockchains aren't a one-size-fits-all technology. Each solution should be specifically constructed to address the complexities of the problem it's intended to solve.

This does not, however, mean that blockchains are destined for use only in fringe scenarios that traditional databases cannot or do not serve. In some cases, existing processes and the use of traditional technology are extraordinarily inefficient, expensive, or unsecure. In other scenarios, blockchains will open up opportunities that were previously impossible.

Having laid out the technical foundations of blockchain technology and its strengths and weaknesses relative to traditional peers in Parts I and II, respectively, this chapter will outline a decision-making process when considering a blockchain deployment. It will seek to clarify the important questions we believe leaders should ask when determining whether a blockchain is well suited to solve a given problem.

Given the state of the technology today, we believe that blockchains are an appropriate solution to consider only when multiple parties that don't trust each other or can't coordinate and cannot agree upon a suitable or cost-effective intermediary (or group of intermediaries) want to transact or share and reconcile data.

Further, we recommend that every potential implementation be framed in the context of a case-specific cost/benefit analysis to ensure that theory does not outpace practicality. Only after determining that distributed ledger technology makes theoretical and economic sense should organizations try to decide what type of blockchain or other type of DLT is best suited to their needs. Each type has its own benefits and drawbacks, and the decision should be made in consultation with a technical partner and other stakeholders.

The structure of most of this section will follow that of the decision tree located at the end of the chapter and should be viewed in conjunction with it. Each subsection will elaborate on certain questions in the tree. The chapter will then close with a list of principles for deployment that we believe are important regardless of the use case and type of blockchain being built.

As with our evaluation of blockchain technology's strengths and weaknesses in Part II, we make the recommendations in this chapter based on our understanding of the technology as it stands today and our expectations of how it will develop in the next few years. Considering the speed at which the space is evolving, we recognize that our forecasts may prove conservative.

Is there a problem with your existing solution?

The enthusiasm surrounding blockchains is exciting for those of us who believe in their potential, but indiscriminately applying the technology for no reason won't help anyone. Therefore, organizations should start by asking a couple basic questions before considering a blockchain:

- Is there an issue with the existing solution to your problem? If the solution currently in place is able to serve users securely and efficiently, it is unlikely that implementing a blockchain would be cost effective at this point in the technology's evolution.
- Are there any alternative new technologies to better solve your problems? Blockchains aren't the only novel technology that may be able to address problems with a legacy solution; organizations should consider multiple angles of attack before deciding to pursue a blockchain.

While traditional databases or other, more mature, technologies will be sufficient or superior to blockchains in some applications, there are cases in which the legacy solution has proven inefficient, unsecure, or expensive. It may simply be the case that no such legacy technology exists. In these scenarios, a blockchain may prove capable of addressing issues with the older solution and merits consideration. We note that the outcomes of early experiments with the Internet in the 1990s are cases in point of the advantages that can arise from experimentation when such testing can be done in a close to cost neutral fashion.

Are prerequisites in place?

Blockchain technology cannot function at its full potential without certain prerequisites, which differ depending on use case. There are, however, two that are consistent across deployments: Internet connectivity and digitized data.

As we noted in Part I, distributed consensus relies on exchanging data over a peer-to-peer network. This requires Internet access. As such, Internet infrastructure of some kind is necessary for a blockchain to function. Not all participants in a blockchain network will need the same level of connectivity: the nature of that connection can differ depending on what it is being used for. Full nodes in large public networks like Ethereum will require high speed connections that can handle large amounts of data, whereas some entities, such as "lite nodes" or "internet of things" (IoT) devices, may not require as consistent or fast a connection. As the AgUnity and WFP cases later in this paper show, organizations that operate in suboptimal connectivity environments can get around this hurdle to some degree. That said, such deployments do require some form of Internet connection and are generally limited by the lack of high-quality connectivity.

The second critical prerequisite is a source of high quality digital data. Before a piece of information can be added to the blockchain, it must be in accurate digital form. This is true regardless of the type of chain or where the actual data is being stored (on or off-chain): deriving a dependable digital fingerprint requires that the information being input to the system be accurate. As the Lantmäteriet Case Study will show, in cases where organizations are

operating with pre-existing records or off-chain data, they should first ensure that the information is accurate and digitized.

Improving internet connectivity and digitizing records are expensive and time-consuming processes. They also add enormous value in and of themselves. If all blockchain technology does is provide the impetus to improve global Internet access and to check and digitize physical records, it will have done an incredible amount of good for the world.

Are there "multiple parties" involved?

In cases in which an organization has no suitable existing solution and in which necessary prerequisites are in place, one can proceed to ask whether a blockchain makes sense for the use case in question.

The first question is simple: are there multiple parties involved? The most compelling use cases for blockchain adoption concern multiple uncoordinated parties seeking to interact with a specific series of data.

Throughout our research, we have seen numerous references to the use of blockchains as a means to facilitate trust between multiple "parties" or "entities", but few definitions of what a party or entity is in this context. We believe the appropriate units of analysis during a blockchain evaluation are not organizations, titles, or departments in the traditional sense, but databases. If two entities have different database instances to track the same thing, they should be considered separate. For example: a sales department and customer service department within the same company may have two different databases that track the same customer relationship; in another scenario two companies may have two separate databases to track the production and delivery of the same good. If a database is intended for use by a single entity, then a traditional database is all that is necessary.

Blockchains allow multiple different parties to read, write to, and store copies of an identical database, creating a single record that is automatically reconciled across instances. This reconciliation is highly beneficial, but it is not unique to blockchains. If multiple separate entities are writing to a database and they trust each other to be honest and correct, then a blockchain isn't necessary: the parties could instead consolidate data onto a single traditional database that they can all write to, edit, and read. More conditions must therefore be met before a blockchain can be considered a reasonable solution.

Is there mistrust or lack of coordination?

One of the things that makes a blockchain special is its combination of consensus mechanisms and tamper-resistance. These combine so that information can be validated without the need for a third party and make it difficult for anyone in the network to "change history". Thus, in order to justify the implementation of blockchain technology, there must exist some degree of distrust or dis-coordination between the entities involved.

Though blockchains are clearly applicable in scenarios where parties have misaligned incentives, mistrust need not arise from the suspicion that one party is up to something

nefarious (like secretly modifying the database). Even when incentives are aligned, mistrust can be well founded: it can be as simple as one department not trusting another department in the same company not to make a mistake.

Alternatively, the problem may not be mistrust, but a lack of coordination that arises because it is very difficult for parties to independently identify each other. This problem has plagued humanity for millennia, and many of the most successful companies in history have been those that have facilitated associations between otherwise unconnected people and organizations.

Is there a suitable intermediary?

When there is a lack of trust or coordination, separate entities usually rely an intermediary (like a bank, a broker, or a government) to help coordinate parties or facilitate trust. Organizations may also develop situation-specific alternatives that use low-tech processes or traditional technology (often proprietary software) to enable transactions and data exchange to take place. The use of intermediaries and alternative processes can seem time consuming, but it can also be simpler, more efficient, and cheaper than the design and deployment of a blockchain at this stage in the technology's maturity: the evaluation is necessarily case dependent. When intermediaries prove unsecure, inefficient, or expensive, a blockchain may be a superior solution.

Further, one can't always able to find an intermediary capable of filling the role. In some cases, this is because there is a question as to the neutrality of the intermediary; in others, participants may distrust centralization/authority in general. Certain socio-political landscapes can garner distrust within systems, such as rampant corruption or chronic ineffectiveness. In these cases, blockchain can remove the need for an intermediary while maintaining the integrity of the system.

Sometimes mistrust of authority plays no role in the failure to find a suitable intermediary. The involvement of an intermediary may simply be economically prohibitive or impracticable. For example, it may not make sense for a bank to spend millions of dollars complying with KYC (know your customer) regulations to get poor farmers to open checking accounts or to facilitate micro-loans. This problem has led to the large population of underserved, unbanked, and unidentified individuals. By more cheaply and efficiently serving vulnerable populations, blockchain could allow for meaningfully improved outcomes for these marginalized populations.

What is the outcome of a cost-benefit analysis?

As noted above, even if a problem does have an existing solution, there are situations where blockchains have the potential to meaningfully reduce costs and inefficiency. Whether this is the case depends on the existing solution, the maturity of blockchain technology at the time of evaluation, the application in question, and the organization itself (among other things). Consequently, before an organization uses difficulties with an existing process as the justification for a blockchain implementation, it should do an extensive cost-benefit analysis.

While we recognize that not all benefits are quantifiable and that internal hurdles for investment will differ, we strongly advise all organizations to consider these questions before making any major expenditures:

- About how much would it cost to design, deploy, and maintain a blockchain? How does that compare with the cost of upgrading existing systems? (Organizations can ask for rough estimates from technical partners)
- How do the projected cost savings and value-creation arising from the use of a blockchain compare to the expected cost its design, deployment, and maintenance?
- How long would the payback period on the investment in a blockchain be? What is the theoretical return on investment? How does that compare to alternative investment opportunities?

An evaluation of the costs and benefits of a blockchain shouldn't be restricted only to situations where the novel technology is being compared to an inefficient or costly legacy solution. An equally thorough appraisal should be made when confronting a problem for which there is no solution at all: blockchain technology is so early in its development that the costs of design, deployment, and maintenance of a blockchain today may exceed the potential benefits. The calculus will be case specific, but it should be confronted honestly. As painful as inaction may be, the widespread deployment of unjustifiable blockchains is likely to hinder the progress of the technology both within organizations and across industries.

HEXA
FOUNDATION

Identify stakeholders and key participants

Once an organization has determined that it is trying to address a problem involving:

Data submitted by multiple parties in which those parties
Either mistrust each other or cannot coordinate and where
A suitable intermediary can't be found or where
DLT may be able to significantly improve the cost, security, or efficiency of an existing process, a type of distributed ledger could be a suitable solution. In that case, the organization in question should determine which type of DLT or blockchain is the best fit.

Due to the complexity of designing and developing a scalable blockchain solution and how hard it is to change a blockchain's protocols if an error is made, there are few scenarios where we believe it is prudent for all but the most technically-savvy and well-funded of organizations to attempt to make these decisions on their own.

Any blockchain solution will exist in the context of established systems and structures. As a result, for a blockchain to meet its stated goal, it is critically important that it meet the needs of potential users while aligning with what is legal and technically feasible. Therefore, though we provide further guidance as to what type of solution may be appropriate, choosing between different types of blockchains and DLT should not be done alone: technical partners, legal and regulatory experts, and (most importantly) the system's potential users and stakeholders should be consulted throughout the process of devising a protocol.

Some of the key questions that should be discussed with these parties have been well enumerated by the Beek Center at Georgetown University:

How will the governance structure of the blockchain be created and maintained?
How will you establish identity in the network?
How will you define and grant access and privileges to network members?
How will you verify and authenticate data and transactions?
How will you define and grant ownership of the data that resides the network?
How will you secure your network?

A thorough understanding of the qualitative answers to these questions is important to have, regardless of the degree to which an organization or coalition is outsourcing the design and development of its blockchain.

How important is it to identify network participants?

In order to determine what kind of blockchain or other DLT is best suited for a particular solution, the first question to ask is whether network members care about who is validating a transaction or participating in the network. Among the reasons why one would want parties to be identifiable are:

- One's organization is in a heavily regulated industry, like finance or healthcare
- One wants to be able to hold parties accountable for their actions in off-chain settings, such as courts
- One wants to make sure that only parties with large stakes in the success of the blockchain are authorized to validate transactions
- Data used as part of the validation and/or verification process arises from off-chain transactions or events
    - Data regarding on-chain events can be validated according to the blockchain's protocols alone. Trusting off-chain data requires a belief in a validator's qualifications and reputation. This in turn, relies on knowing aspects of the validator's identity.

If being able to identify members of the network is important, then an organization should choose between a private blockchain and a distributed ledger: as of the time of writing, public chains are open platforms that utilize economic incentives to drive good behavior and don't require nodes to provide any identification. Private chains and distributed ledgers, on the other hand, generally require identification of some sort to be allowed into the network. We note that if decentralized identity solutions are successfully developed and adopted, this differentiation between public and private chains may no longer be as relevant.

How sensitive is on-chain information?

The next question to ask is: does it matter who reads the contents of the chain? Among the reasons why privacy may be an issue are:

- There will be actual data stored on the blockchain (as opposed just digital fingerprints derived from data stored off-chain)
- Information on the chain, even metadata, is highly sensitive
- There are legally mandated privacy requirements

Some organizations may want to restrict the legibility of information to a select number of entities. If that is the case, they will have to choose between a private chain and other types of distributed ledger. However, in certain cases the data being exchanged (and the associated metadata) may be so sensitive that a blockchain isn't an appropriate solution: as we noted in Parts II and III, a blockchain is typically more transparent than a traditional database is. While certain blockchains can be structured to maximize anonymity, in instances where even the metadata surrounding a transaction is sensitive, we would recommend organizations explore an alternative distributed ledger technology. Some alternative DLTs are less transparent than

even private blockchains are but share some of a blockchain's other characteristics. Depending on the application, they may prove suitable.

Does your blockchain deal with off-chain data?

If validator identity and read-permissions aren't an issue, a public blockchain may be a suitable solution. To make the decision between a public and private blockchain at this point, we recommend organizations first ask whether they are engaging exclusively with on-chain data (like that arising from the exchange of cryptoassets). As noted in Part II, blockchains do not have the ability to natively validate data that doesn't result from on-chain interactions. In those cases the application of blockchain technology should be approached with caution.

Where organizations are dealing only with on-chain information, validation is driven entirely by code, and both private and public blockchains may be suitable options with individual puts and takes that must be considered. In the decision tree at the end of this chapter, we indicate a preference for public chains due to their transparency, openness, and (theoretically) greater security.

If the data in question involves off-chain assets, organizations must also assess whether those assets can be accurately tracked and represented on the blockchain by means of tokenization or a similar process.The easiest assets to manage and track using blockchains are those that are fungible, like commodities and identical digital files, or those that have characteristics that are very difficult to change/imitate, like diamonds and unique digital assets. The more easily that an item can be changed, copied, or differ in quality (and have that change matter to a buyer), the more difficult we believe the asset will be to effectively deal with using a blockchain. That's not to say that blockchains can't theoretically be used to manage off-chain assets that don't fit into either of the above categories.
If an asset or data cannot be fully expressed in a digital format, organizations will require the use of additional tools, procedures, rules, and enforcement structures to ensure that analog data that is processed off-chain can be captured and recorded accurately on the blockchain.
As the technology matures, we expect that the number of asset classes that can't technically be managed on chain will decline; however, in scenarios where an asset cannot be easily tokenized and tracked, blockchains should be, at most, used in a redundant capacity.

In the event that an off-chain asset can be effectively tokenized and tracked, we suggest organizations use a private blockchain; in our opinion, the infrastructure necessary to rely on public blockchains as a means to exchange off-chain assets is not yet in place. This conclusion may change if the public chain ecosystem develops improved ancillary infrastructure, decentralized identity solutions, and reputation systems; for the moment, however, code is not law.

Closing

Like the Internet, neither blockchains nor their DLT brethren are a panacea; as such, the decision to implement them should not be taken lightly. We hope that this section will have provided structure to our readers' decision-making processes and will lead to a more measured approach to blockchain technology in general. While blockchains may help to

facilitate trust, the sales process surrounding them and their associated infrastructure is no different than that of their traditional database peers: promotional, at times to the point of inaccuracy and fraud. In all cases, buyer beware.

HEXA
FOUNDATION

Do you need a **blockchain?**

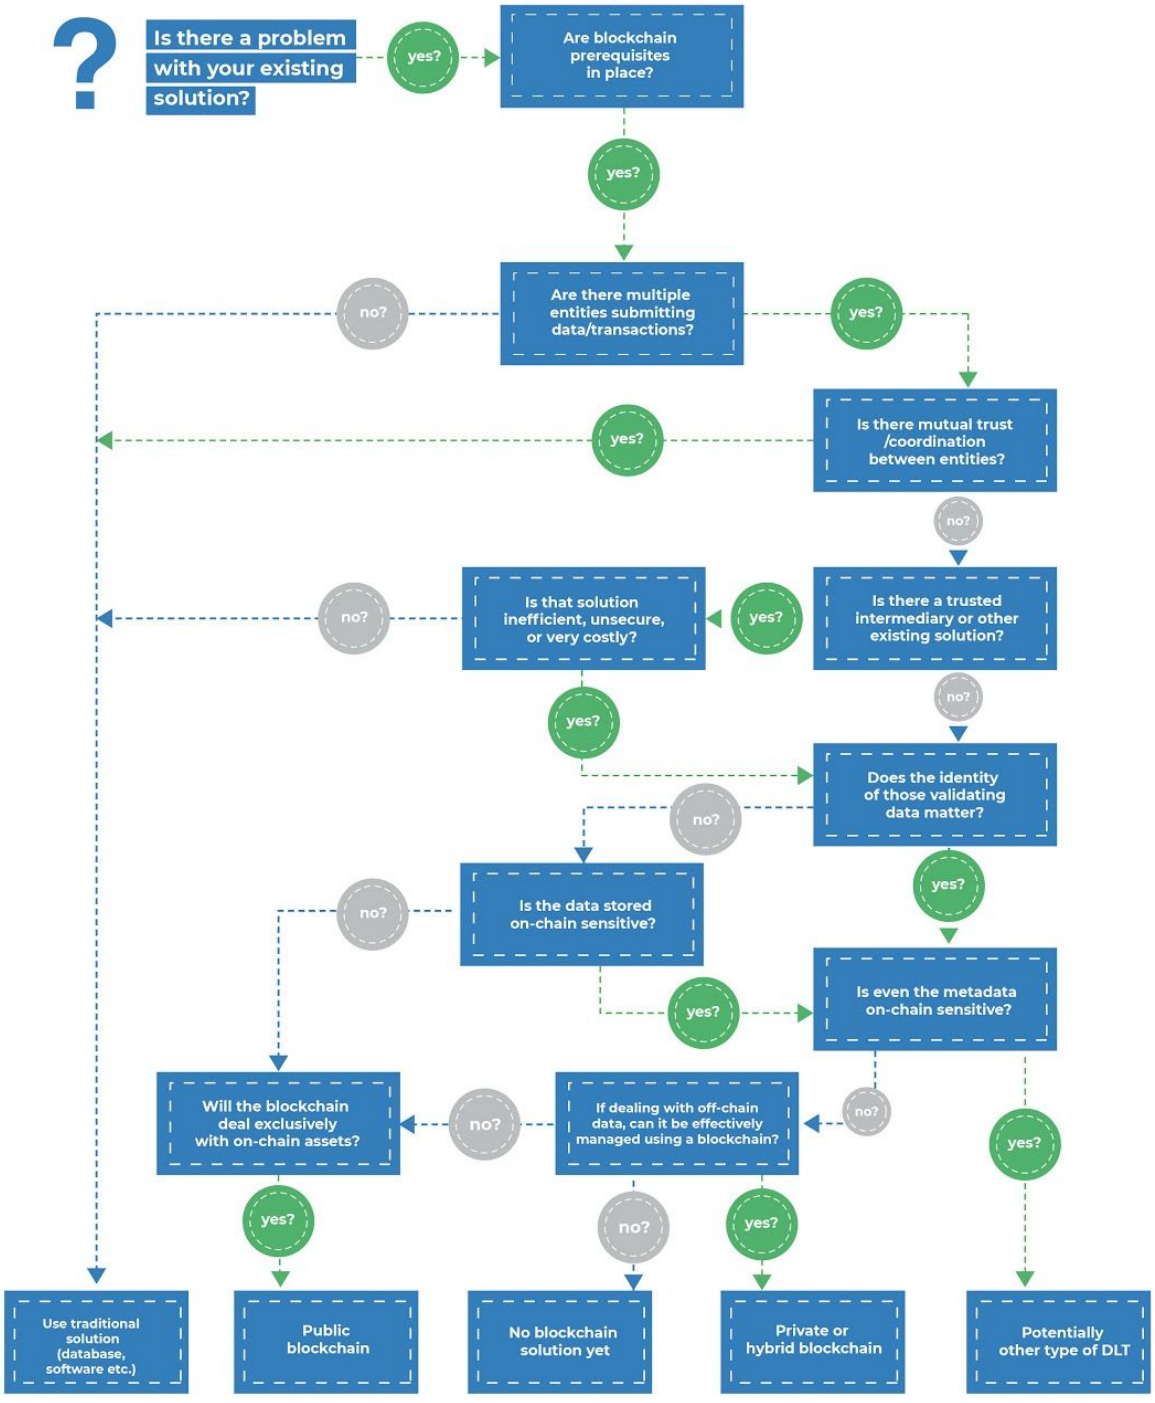

HEXA FOUNDATION

Part IV: Key Principles for Blockchain Protocol Design and Deployment

Once an organization has defined its problem and goals, narrowed down the type of blockchain or DLT it wants to design, and checked that a blockchain makes economic and technical sense, coalitions can begin to consider the details of the design and deployment. Like those related to cost/benefit analysis and design, design and deployment decisions will be case specific; however, our research has led us to develop a number of principles that we believe are broadly applicable. As some readers will note, much of what we include herein is not applicable solely to blockchains: many of these principles are best practice for any novel technology. We include them here because we consider them an essential part of the process of deployment and design.

1st Principle: Ensure the presence of prerequisites

As we noted in Part III, high quality internet connectivity and accurate digitized data are necessary if a blockchain or distributed ledger is to function as intended. Though not all participants in a blockchain network will need the same level of connectivity, bandwidth and connection quality are limiting factors that designers should consider before beginning to plan or deploy a solution. Similarly, accurate digital data is essential in cases where previous records are relevant both to ensure that the solution accomplishes its stated goal and to avoid the blockchain becoming a permanent repository of erroneous data. Organizations considering a blockchain solution should therefore ensure that if off-chain or historical information will be a part of the system, that those data points are both accurate and in a digital form that can be used by network participants.

2nd Principle: Consult with potential stakeholders and participants

Though blockchains are often viewed as purely technical solutions, they are only valuable if they benefit from a large network; consequently, building a successful protocol has an important sociological component that is hard to code. A blockchain network without participants is worthless, so making sure that the chain is designed to serve the purpose of potential network members is very important.

Consequently, maintaining communication with constituents is important throughout a blockchain's deployment process. This can be difficult to accomplish, especially when trying to help underserved populations. While it is unreasonable to expect organizations to deploy blockchains in a way that satisfies everyone, talking to partners users throughout the process will ensure that the solution is addressing pain points and avoiding impasses. The result will be a smoother rollout and higher probability of success.

3rd Principle: Design blockchains like hardware, not software

Though blockchains are technically a type of software, organizations should recall that unlike that of their traditional peers, blockchain governance is decentralized. This means that making changes to a protocol (even to eliminate coding errors) typically requires gaining the agreement of at least a majority of the chain's voting stakeholders.

Because changing a blockchain requires the agreement of so many more parties than changing a traditional software program does, the technology is currently poorly suited for the traditional software deployment process. Blockchain deployment should more closely resemble the development of hardware components, which undergo extensive testing before being produced at scale.

This is not to say that experimentation isn't possible or encouraged after a solution is launched: later in this section we recommend the use of pilots precisely because they allow organizations to identify problems with a protocol before it becomes too big to modify.

4th Principle: Maximize interoperability and avoid vendor lock-in to the degree possible

Blockchain technology is immature, widely accepted design standards are not yet in place, and the sustainability of today's major protocols and platforms is uncertain. We therefore suggest that organizations maintain flexibility during the design and deployment process by ensuring the maximal degree of interoperability (between platforms) and avoiding vendor or platform lock-in, if possible.

A number of companies are working toward facilitating interoperability between distinct protocols, but it is still a matter of concern that should be addressed with technical advisors. Avoiding lock-in is a separate challenge; a starting point is trying to ensure that data can be moved from one platform to another and avoiding blockchains with strange data formats, exclusive coding languages, or that require storage of all data to be on that platform.

5th Principle: Start small and run redundantly

Given the complexity of designing and implementing a blockchain and the costs that can arise when trying to scale it, we strongly recommend that all organizations begin with a proof of concept (PoC) followed by a pilot before proceeding to a full roll-out. These provide a compartmentalized environment that will allow for effective measurement of results and evaluation of the cost-effectiveness of the solution on a small scale, both in absolute terms and relative to the legacy solution (if there is one). Staging a PoC and pilot will also allow for easier (but still not simple) modification of the protocol in the event that issues arise and provide an opportunity for experimentation.

6th Principle: Pre-set numerical goals and qualitative signposts for what "success" looks like

In Part III, we elaborated on the importance of ensuring that a blockchain is cost-effective. We also recognize that when implementing a novel technology, it can be difficult to view the results dispassionately. As such, we recommend that before any sort of deployment (even a pilot) has taken place, organizations set both clear quantifiable goals and explicit qualitative signposts that will serve as sanity-checks before making further investments in a blockchain deployment. The results of each step of the process should be compared to what was expected at the start.

7th Principle: Scale slowly and run redundantly

The temptation upon achievement of satisfactory quantitative and qualitative results in early deployments may be to expand the use of the technology quickly. This is generally not advisable. Like any technology, blockchain deployments may experience "growing pains" that won't necessarily have to do with the protocol itself, but with the development of relationships with constituent parties, the growth of an ecosystem etc. These will take time to resolve.

Further, success in a small-scale pilot is, given the immaturity of the technology, insufficient evidence of the sustainability of a protocol. Thus, except in cases where there is no legacy solution, we suggest running the blockchain redundantly with existing systems and to slowly transition more responsibilities from the traditional platform to the blockchain. Do not "rip and replace".

Closing

In this section we have offered a number of principles for design and deployment once an organization has determined that a blockchain may be an appropriate solution for a given problem.

As we have denoted throughout the paper, we expect that as the blockchain ecosystem evolves, the number of situations in which the technology will be both applicable and cost-effective will grow until blockchains are essentially taken for granted and papers like this one are irrelevant. It is at that point that blockchains and their DLT peers will have begun to reach their potential.

HEXA
FOUNDATION

Part V – Case Studies

As we noted in Part III, it is difficult to prescribe specific best practices for blockchain-related design and deployment because use cases can vary significantly depending on the problem being addressed, the network's participants, and existing infrastructure (among other factors). We have therefore performed four full-length case studies of blockchain and DLT deployments in a variety of contexts and believe that they illustrate aspects of the value that the technology can add in the public, private, and non-profit sectors:

- AgUnity's solution, which shows how DLT can be used to increase the income of smallholder farmers in the developing world by 300% through the facilitation of trust in cooperatives
- A Walmart-led consortium's effort to use a blockchain to track millions of grocery items across the companies' supply chain and reduce trace time from 6.8 days to 2.2 seconds
- The UN World Food Programme's deployment, which has helped reduce the cost of cash based transfers to refugees in its pilot by 98% and will facilitate inter-agency coordination
- A consortium led by the Swedish Lantmäteriet, showcasing blockchain technology's potential to help reduce the number of deficient land titles and cut real estate transaction times from 3-6 months to a few days

While the teams at each of the aforementioned organizations were extraordinarily helpful, time and resource constraints have limited the amount of diligence we could do on any one of them, especially given that like the technology, the ventures herein are evolving. Therefore, these studies shouldn't be viewed as an endorsement of any of the solutions described.

We selected these particular cases because they illustrate different ways that blockchains and their DLT peers can create value and because they show how deployments have been managed by different types of organizations; however, as we will show, none of them are perfect. The solutions are all less than three years old, and we are so early in the development of blockchain technology and its associated ecosystem that it would be unreasonable to expect perfection at this stage.

If nothing else, these cases all show that designing and deploying a blockchain solution is complex and can require considerable amounts of time and resources. We hope that the studies, combined with the rest of the paper will provide inspiration and a guide for those organizations considering the technology.

HEXA FOUNDATION

AgUnity Farmer Cooperatives Case Study[17]

A special thanks to Nick Miller, David Davies, and the rest of the AgUnity team for their participation in and help with this case study.

What is the problem being solved?

Over 1.5 billion people around the world are members of smallholder farm households with less than 2 hectares of land[18]. A majority of smallholder families live in abject poverty; many earn less than $3,000 a year in gross income[19]. This works out to under $2 a day per person in a typical five-person family. Ironically, many smallholders spend a significant proportion (over 80%, in some cases) of their income on food, leaving little for durable goods, education, or health-related expenditures[20]. This dynamic creates a vicious cycle: no savings or access to financing makes long-term investment in capital (human or physical) difficult, thereby curbing economic mobility and the long-term development of agrarian economies as a whole.

The UN's Food and Agriculture Organization data shows that farmers typically receive a small fraction of the value that wholesalers earn from the sale of farmers' crops[21]. This differential is due in part to the highly fragmented nature of smallholder farms: over 80% of the world's over 570 million smallholder farmers own only about 12% of the world's farmland[22]. Geographic dispersion, lack of resources, and, at times, illiteracy, make it difficult to coordinate. Consequently, smallholders' negotiating leverage with suppliers and buyers is minimal.

The creation of smallholder cooperatives has been shown to be a significant contributor to higher incomes and improved quality of life[23]. Effective co-ops allow small farmers to demand better prices from both buyers and suppliers and to share best practices and equipment with each other, leading to better results for all members of the group.

Unfortunately, the legacy systems and processes used to create and manage cooperatives are unreliable in less-developed nations. The "lack of transparency, restricted access to price data, lying, graft, and corruption[24]" that result from current practices combine with illiteracy among some smallholders to create a lack of trust in the cooperative system. This mistrust can make co-ops hard to form and maintain. In some cases farmers will end up watching a portion of their harvest rot because they can't find a buyer that they trust. These dynamics lead community members to miss out on considerable value capture.

---

[17] An abridged version of this case study is available in the paper published by the Blockchain Trust Accelerator

[18] Sarah K. Lowder, Jakob Skoet and Terri Raney, "The Number, Size, and Distribution of Farms, Smallholder Farms, and Family Farms Worldwide," *World Development* 87 (November 1, 2016), 16-29. doi:10.1016/j.worlddev.2015.10.041. http://www.sciencedirect.com/science/article/pii/S0305750X15002703. 24.

[19] George Rapsomanikis, *The Economic Lives of Smallholder Farmers: An Analysis Based on Household Data from Nine Countries* (Rome: The Food and Agriculture Organization of the United Nations,[2015]). 1 and 21.

[20] Ibid, 26.

[21] "Prices,"accessed Apr 17, 2018, http://www.fao.org/prices/en/.

[22] Lowder, Skoet and Raney, "The Number, Size, and Distribution of Farms, Smallholder Farms, and Family Farms Worldwide,"16-29. 1.

[23] Sripad Motiram and Vamsi Vakulabharanam, "Corporate and Cooperative Solutions for the Agrarian Crisis in Developing Countries," *Review of Radical Political Economics* 39, no. 3 (September 1, 2007), 360-367. doi:10.1177/0486613407305284. https://doi.org/10.1177/0486613407305284.

[24] "AgUnity: Blockchain for the Greater Good," accessed April 17, 2018, http://www.agunity.com.

While the AgUnity solution is still a work in progress, it serves as an example of distributed ledger technology being used help underserved populations. The case exemplifies some of the difficulties nonprofits and social-impact ventures face when trying to deploy DLT to serve extremely needy populations. It is also an example that demonstrates the importance of understanding network constituents and establishing partnerships when designing and deploying a solution.

The story

AgUnity was founded in 2016 as a for-profit enterprise to design a simple DLT-based solution to facilitate trust in smallholder farming cooperatives, thereby improving resource pooling, communication, and the lives of the members of AgUnity Networks. The organization's structure is still evolving, but its goal continues to be to increase farmers' income through the expansion of the cooperative model and then use the platform and the cooperative's scale to allow members to access goods and services like solar lamps, mosquito nets, microfinance offerings, and crop insurance at affordable prices.

AgUnity considered traditional alternatives to distributed ledger technology as a means to achieve their goal, but determined that a traditional database and software solution wouldn't offer the capabilities they sought. Namely, they found in their conversations that farmers no longer trusted middlemen at all: they wanted a permanent record. The founders of AgUnity felt that only DLT could offer the transparency, security, and tamper-resistance necessary to facilitate trust among the company's users.

Technical points of the solution

The company's product is a combination of a hardware and software solution: each farmer is provided with a low-cost smartphone that has been wiped and re-imprinted with a proprietary operating environment and application designed with a particular crop in mind (wheat, cocoa, rice, etc.). The rationale for using devices rather than just an app was threefold: first, many of AgUnity's target demographic don't have smartphones or have phones that vary in quality and security. Second, the device combines with a passcode to serve as a means of unique identification for its owner. Third, it allows AgUnity to imprint its own operating system, which helps the AgUnity app run smoothly, makes support easier, and allows for the transfer of identities to a new device in the event that the old one breaks or is lost/stolen. The phones are otherwise normally functional.

The primary component of the AgUnity operating system is a distributed ledger application that is currently built on an iteration of the Multichain blockchain platform. We note that AgUnity is considering a move to an alternate platform, and will maintain flexibility to switch again as necessary.

Each smallholder community has different needs and desires, and AgUnity's solutions are designed with specific crop types and cultures in mind rather than applying a "one size fits all" model. The largely pictographic application helps farmers enter into agreements with representatives of their local co-op and create a permanent record of the terms (mostly what

volume of crops has been given to the co-op) that is validated by both co-op representatives and the farmers themselves.

Because Internet connectivity can be spotty in many of the less developed countries in which AgUnity operates, the app is designed to allow for off-line transactions that are posted to the ledger once connectivity is re-established. The phones are not nodes, but copies of the ledger are stored on multiple nodes belonging to AgUnity and to a number of NGOs. Anyone can theoretically have an AgUnity node and act as an independent record-keeper.

The application also provides a near-real-time record of crop processing as cooperative representatives post the completion of each step (for example, fermenting and drying cocoa beans) to the ledger to maintain transparency and trust. There is currently no verification process in place, but AgUnity claims that the fact that there is a record at all reduces the risk of cheating much like the threat of a tax audit does.

Once the entire co-op's harvest has been sold, the amount earned is posted to the ledger for all to see and a record of each farmer's share of the proceeds is automatically posted to their wallet. The farmer then goes to the cooperative to collect his share of the proceeds in cash. AgUnity's solution provides a source of truth (a receipt) for the farmers, but does not currently use smart contracts or other means to automatically execute value transfers. Enforcement of the contract relies on traditional means (like the legal system).

The hardware platform and operating system also allows for the provision of other, non-blockchain related, functionality including a messaging app that helps get farmers involved with cooperative activities and facilitates resource pooling, harvest planning, and crop collection. All of the data is encrypted and belongs to the farmer – AgUnity does not have the private keys necessary to see or sell user data.

The AgUnity team decided that, other than the Multichain core, the operating system, applications, and DLT would be built in-house with the aid of a few contractor programmers. The company's co-founders and CTO are all former financial IT executives and have extensive experience with both security and with DLT; however, the decision to design in-house was apparently driven more by the importance of having had face to face interactions with their users (smallholder farmers) than because the team felt they were better technically qualified.

Scaling

After approximately six months of development starting in early 2016, AgUnity launched a proof of concept and pilot in Nanyuki, Kenya lasting about a year. The Nanyuki pilot yielded important lessons, most of all how difficult it was to build and coordinate a network of smallholder farmers without a local co-operative that has existing relationships. A blockchain without a network is useless, and networks in these communities often require local knowledge. Consequently, pilots are now only launched in locations where there is a local NGO partner and/or co-op with existing connections to the community to help distribute phones and manage relationships.

The positive results in the Nanyuki pilot led AgUnity to launch another iteration of the solution on Bougainville Island in Papua New Guinea in 2017 focusing on cocoa instead of wheat. This pilot has yielded extraordinary results. Between the Nanyuki and Bougainville pilots, the solution has been deployed to about 50 users**,** and resulted in a 300% increase in income for the farmers equipped with an AgUnity phone. While some of the improvement may have been due to better harvests (irrespective of AgUnity), the company claims that a significant proportion is due to higher crop volumes sold (due to less of the product rotting) and higher sales prices resulting from co-op membership. With those results in hand, AgUnity is working to expand its Bougainville pilot, and will launch new pilots in multiple other developing countries, including a partnership with the World Food Program in Ethiopia. The goal of these pilots is to prove the viability of the solution in other countries and when dealing with different crop types and local cultures.

The company is also in the process of adding a variety of functions to the application, among them settlement on-chain, which will allow for the transfer of actual funds directly to a farmer's electronic wallet using services like M-Pesa and effectively creating a bank account for the farmer. As exciting, is the opportunity to help farmers to use their harvest to securitize loans. The wallet functionality will allow for the provision of "microliens" (so-called because payments are automatically withdrawn from farmer's wallets), and also permit lenders, NGOs, and other organizations to control what their money is spent on: grants and loans will be transferred in digital form to the farmer's wallet (where their use can be restricted) rather than provided in cash.

Areas for Improvement

The AgUnity solution is still in very early stages of development, and we have identified a number of areas for improvement. First, the solution is somewhat centralized. Though the company can't change the ledger, AgUnity is in partial control of user identities because it is the only entity that has the "backup keys" necessary to move the digital identity of farmers to new devices. Further, in order to help the application function in no-connectivity environments, the transactions are first posted to a cloud database and then to the chain. AgUnity could, in theory, try to modify a transaction record before it is posted to the chain if the company were able to ascertain private keys for both the cooperative representative and the farmer in question.

Second, while the ledger acts as a record for a transaction, there is still a level of trust involved considering that there is no independent validator overseeing the processing and sale of the product, and farmers still rely on off-chain enforcement in the event that a cooperative defaults on an agreement. This is suboptimal given that some locations may have legal systems biased against the farmer, and exemplifies the fact that blockchains are not yet ready to manage transactions involving off-chain assets on their own.

Finally, while we understand the rationale for it, the fact that the solution is currently hardware-based imposes limitations on how big an impact it can have. Because possession of a customized phone is required to join a network, the AgUnity solution will be inaccessible in geographies without an NGO or existing co-op that has a relationship with AgUnity. This significantly limits the application's addressable market and scope. Further, each phone is

programmed for a certain type of crop in a particular location, limiting individual farmers' ability to shift to other types of crops and the ease with which the solution can be applied to new scenarios.

Conclusion

This case study has shown an example of the design, implementation and deployment process for a DLT solution by a small organization trying to tackle a big problem that traditional technology has not been able to solve. As we have shown, there are still a number of areas in this case that require attention from a technical perspective, but the outcomes that AgUnity's application has produced in Nanyuki and Bougainville are promising if they can be replicated.

The company's design, implementation, and deployment process was largely aligned with our recommended principles. The solution was built by a technically skilled team on a pre-existing platform, and AgUnity maintained flexibility and avoided vendor lock-in. While the company did not set numerical goals in its early pilots, the deployment process has been measured and conditional upon satisfactory improvement in farmer income in Nanyuki and Bougainville. Most impressive has been the effort that AgUnity has put into understanding the potential constituent members of its networks before beginning to design and deploy the solution. These efforts have ensured that the company counts with the support of pre-selected on-the-ground partners and that the application is tailored the needs of the community where it is being deployed.

AgUnity's process did not however, conform to all of our suggested principals. Most importantly, the geographies in which AgUnity is deploying typically do not have widespread Internet connectivity. The lack of high quality connectivity in farming communities has forced AgUnity to develop the workaround described above, which is a point of weakness in the technical infrastructure. In this matter, the company has little it can currently do. It has made efforts to secure the database itself and is looking at a number of alternative ways to resolve the connectivity problem.

As the organization shifts from a for-profit corporation to a non-profit entity, AgUnity will continue to refine its solution and deploy it in a 3-5 new locations a year. As the solution is only about two years old, its final form has yet to be determined. Our evaluation has concluded that there is more work to be done before this can be considered a template for other implementations for the underserved, but it is off to a strong start.

HEXA
FOUNDATION

Walmart Supply Chain Case Study[25]

A special thanks to Drew Sadler, Frank Yiannas, Tejas Bhatt and the rest of the Walmart Food Safety team for their participation in and help with this case study.

What is the problem being solved?

Humanity consumes approximately 5,133 metric tonnes of food every minute[26]. The path that food takes from "farm to table" is extremely complex: not so much a linear supply chain as a "food system" dependent on hundreds of thousands of different entities around the world. This system provides consumers, especially in developed nations, with a more convenient means to attain a broader selection of cheaper food than ever before. It also, however, presents its own set of challenges, especially for the retailers and brands that the public trusts to provide them with safe, nutritious, and accurately described choices.

Unfortunately, collating and analyzing information regarding the source of food and the path it takes to a retailer's shelves is an extremely labor intensive and time-consuming process. Many of the participants involved in the production and distribution of food still use paper-based systems to manage records. Even if information is captured in digital form, the data is usually siloed in disparate systems that are incapable of interoperation or even communication with each other.

The inefficiency of the current process imposes significant costs on both consumers and the companies that serve them. Consumers can be exposed to unsafe or inaccurately labeled foods and also end up absorbing some of the expense associated with the inefficient process of managing retailers' food systems. Retailers and brands are forced to undergo costly and redundant processes to try to trace the source of their products and are held accountable for the outcomes of events that are often beyond their control (like outbreaks of food-based illness).

Despite attempts by many brands and retailers to solve this problem using traditional technology, the food system experiences adverse events time and again: in 2006, the United States was hit with a nationwide outbreak of E-coli linked to bagged spinach that led retailers and restaurants to pull all bagged spinach, regardless of source, off of store shelves and menus. It took two weeks to trace the source of the outbreak. Twelve years later, the US is facing an almost identical issue. On April 10th 2018, the FDA reported an outbreak of E-Coli linked to romaine lettuce[27]. Though regulators narrowed down the source of the outbreak to Yuma, AZ, they have still failed to pinpoint the exact origin of the tainted produce. The price of these types of outbreaks is higher than the millions of dollars in losses to retailers, brands, and farmers; the April 2018 outbreak went on to kill five people and hospitalize 96 more across 36 states. As this case study will show, blockchains may offer a better way forward.

---

[25] An abridged version of this case study is available in the paper published by the Blockchain Trust Accelerator
[26] "World Food Clock,"http://www.worldfoodclock.com.
[27] "Multistate Outbreak of E. Coli O157:H7 Infections Linked to Romaine Lettuce," last modified Apr. 10, accessed May 5, 2018, https://www.cdc.gov/ecoli/2018/o157h7-04-18/index.html.

This case will focus on how Walmart's use of blockchain technology to track the food sold in its stores has evolved into an effort by ten of the world's largest retailers and brands to improve accountability and safety across the global food system. It will show how the technology has facilitated cooperation among many, sometimes competing, entities to successfully track millions of food packages across supply chains belonging to companies including: Walmart, Kroger, Wegmans, Tyson, Driscolls, Nestle, Unilever, Danone, McCormick, and Dole.

The story

The Food Safety Group at Walmart evaluated a number of alternatives as the company and its customers were impacted by food quality events every year. The way the traditional system works is that each participant in the path a particular item takes from the farm or fishery to the shelf has to track food "one step forward and one step back." Retailers then have to piece information together, if they can at all: much of the data is only available to regulators in the event of an outbreak.

At the end of 2016, Walmart's Food Safety group began to consider blockchain as a means to enhance the traceability of the food sold across the company's 12,000 stores and websites. The team found the technology had many of the qualities they were searching for. The tamper-resistant, decentralized, and relatively democratic nature of blockchains was unique and would be critical from both a functionality perspective and the perspective of incentivizing the company's partners (farmers, food processors, brands etc.) to participate. Further, while not unique to the blockchain, the introduction of a novel solution could provide an impetus for all members of the company's supply chain to use standard data formats and a single system, significantly simplifying the tracking process.

Walmart's pitch to its suppliers was simple: using Walmart's blockchain solution would reduce their costs and improve operational efficiency. The technology would allow them to partake in a smaller number of redundant processes, suffer from a lower incidence of unnecessary recalls, and gain insights into the path their products took to the shelves. Though the company didn't comment, we suspect that the corporation's scale was also factor in driving adoption: it is hard for any supplier to say "no" to Walmart.

Technical points of the solution

After evaluating a number of options in a competitive tender process, the company settled on a version of IBM Hyperledger's "Food Trust" technology. The solution is a private blockchain with a permissioning layer on top of it so that the legibility of certain information can be further restricted (in this respect it's a lot like a distributed ledger). Most of the data itself is stored in a traditional database and a record of it is posted to the chain. These technical choices were driven by the desire to align with nascent industry standards, address the importance of identifying members of the supply chain, maintain high scalability, and provide assurance to supply chain partners that their data would remain private. The company continues to explore means to enhance transparency while also ensuring that sensitive information remains confidential; however, the solution currently prioritizes discretion.

Each stakeholder in the supply chain is responsible for uploading data regarding any physical transactions they are a part of to the system. The data they upload usually includes the “who, what, when, and how” regarding an item’s path. Participants are encouraged to use standards like GS1 Global Location Numbers, Global Trade Item Numbers, and a particular type of barcode label to maintain consistency of inputs across the chain. The blockchain then links those inputs together to form a complete picture of the path that food takes from field or fishery to the shelves. Verification of inputs is achieved by comparing the inputs from the party handing off the product and those uploaded by the party receiving it. Walmart also does data analytics on the contents of the chain based mostly on common sense checks: for example, if the product was received before the supplier said it was shipped, it raises a red flag. The Walmart team noted that because erroneous inputs are immutable, the blockchain incents participants to make sure that the data they input is accurate the first time. It creates a sense of responsibility and accountability that doesn’t exist in a traditional system.

Scaling

Walmart began by piloting its solution in its pork and mango supply chains in the spring of 2017. The company chose those supply chains because they were relatively straightforward but also quite different from each other, allowing the solution to be tested in different scenarios. The results of the pilots were exceptional: time to trace products from shelf back to the farm was reduced from 6.8 days to 2.2 seconds. Further, the data from the trials led Walmart to determine that use of the blockchain would meaningfully reduce costs, thereby allowing the corporation to be more competitive and pass some of the savings onto its customers.

After the success of the initial pilots, the company recognized that expanding the blockchain network beyond Walmart could create additional value. As a result, Walmart and Hyperledger approached nine other large companies in the summer of 2017 and formed a consortium to explore collaboration around the technology. As of August 2017, a solution tracking items in the Walmart, Kroger, Wegman’s, Tyson, Driscoll’s, Nestle, Unilever, Danone, McCormick, and Dole supply chains has been live.

Despite the fact that several of the consortium partners are competitors, the companies involved determined that a more open and collaborative solution would be more valuable. It would benefit from a larger and more data-rich network, and avoid the data siloing that plagues the traditional system. Further, it would be cheaper to distribute the costs of expanding and operating the system across multiple entities than for all of them to create solutions in isolation. Walmart convinced competitors by noting that collaboration in this regard would increase the odds of success, whereas the creation of competing solutions would place unreasonable burdens on suppliers and likely hamper each of their individual efforts.

The consortium's solution is currently being used to track about 20 SKUs (a retail term to describe a specific type of item) and has recorded hundreds of thousands of traceability events. Despite the fact that the solution has already tracked millions of items around the world, 20 SKUs is a tiny fraction of the total number of food items sold by Walmart and its partners. The plan for the immediate future is to scale within and beyond the initial ten companies involved and into food sectors that are more ready than others for participation. These sectors are likely to be those that already have certain prerequisites in place, namely standards around identification and labeling of items and their location.

The first of these sectors are leafy greens providers. By the end of January 2019, all direct suppliers who provide leafy greens to Walmart will be required to be a part of the blockchain, and by the end of September, all suppliers will be expected to have integrated their vertical systems end-to-end onto the platform.

The blockchain has intentionally been designed to have a low barrier to entry to ensure that all members of the supply chain, regardless of size or sector, can partake in and benefit from its use. Walmart expects that eventually all of its food suppliers will share data via the solution; however, the company is being intentional about how it expands the deployment in order to ensure that all participants are deriving value from it. Despite the initial success of the blockchain system, Walmart and its partners all continue to use legacy systems as a redundant means of tracking food as an added measure of security and will do so for the foreseeable future.

As the volume of data in the system grows, the options available to Walmart and its partners are tantalizing. In future, Walmart may use the system to optimize the company's supply chain, reduce food waste, and enable consumers to have greater transparency into the source of their food and how it was produced. The hope is that these additions will result in a safer, more efficient, and sustainable food system.

Areas for Improvement

As with all the other solutions explored in this section of the paper, the consortium's solution is still relatively immature and has some shortcomings. From our perspective, the most important of these is that it still relies to some degree on trust above and beyond that which is typical of a private chain due to the fact that it involves physical assets. As we noted in Part III, the blockchain ecosystem has not yet developed a means to reliably validate data resulting from off-chain interactions, especially those involving physical things. Therefore, the Walmart solution still relies on supply chain participants to input accurate information. The reliance on trust is especially important given that the state in which food has been kept and its quality are currently hard to audit and matter a lot to Walmart and its partners. As a result, even if data on the path goods take is accurate, there is no assurance that information regarding their quality is.

Walmart and its partners are aware of this shortcoming, and the company noted that the solution is intended not to be trustless but to encourage a "trust but verify" mentality. As such the consortium is exploring innovations using artificial intelligence and IoT devices to help audit and verify data regarding quality and other inputs, thereby improving the solution.

Conclusion

This case study has shown how a blockchain deployment by a single company evolved into an effort by a group of large retailers and brands to improve the efficiency and safety of their global food supply system.
Supply chains are an archetypal example of a scenario in which a blockchain can add value. They involve multiple different parties that don't entirely trust each other and auditing them currently relies on a process that is both highly inefficient and expensive.

The progression of Walmart's current solution has generally aligned with our principles for deployment. As we have shown, the implementation was proposed with the appropriate prerequisites in place, designed carefully and with the input of partners, scaled slowly only after achieving success in early pilots, and run redundantly with the legacy solution. We note that even Walmart, a company with billions of dollars in revenue, didn't attempt to build its own solution from scratch.

One area where the solution did not conform to our principles is that though the company emphasizes the importance of interoperability and continues to work with the blockchain community to maximize it, the risk of vendor lock-in (to Hyperledger) seems to be quite high: porting data from Hyperledger to an alternative is likely to be difficult when so many entities are involved. This may be a reasonable price to pay for attaining scale.

Despite this deviation and the residual reliance on trust, the solution that Walmart and its partners have built is likely one of the largest deployments of a private blockchain in the world. The deployment is already having positive qualitative effects: Walmart has observed significantly greater collaboration around supply chain improvements among members of the consortium and a higher degree of trust and accountability now that all members of the group as incented to "do the right thing every time", not just when it's profitable. As it grows to include more SKUs and companies, the consortium's chain has the potential to save billions of dollars and enhance the safety and sustainability of the global food system, benefiting consumers and companies alike.

HEXA FOUNDATION

UN World Food Programme Building Blocks Case Study[28]

A special thanks to Jamie Green, Sam Ng, Berhard Kowastch, and the rest of the World Food Programme Team for their participation in and help with this case study.

What is the problem being solved?

At a time when global GDP is at an all time high and when technology has impacted the lives of millions in the developed world, it can be easy to forget that a significant portion of the world's population has been left far behind, and that that portion is growing despite the wealth being created around them.

800 million people experienced food insecurity in 2017[29]; 1 in 113 of people worldwide was forcibly displaced, (up from 1 in 160 a little more than 20 years ago[30]); 22.5 million (roughly a third) of those displaced are refugees fleeing violence, persecution, natural disaster, and famine[31]. Helping these individuals is an overwhelming task that is left to organizations like the UN World Food Programme (WFP).

The Azraq and Zaatari Refugee Camps in Jordan exemplify many of the challenges organizations like the WFP face when trying to care for those in desperate need. The camps opened in 2014 and 2012 respectively and now house over 100,000 people who are victims of the civil war in Syria[32]. Some suffer from serious physical or psychological injury. Many lack formal identification. All lack significant resources.

The challenges the WFP faces are exacerbated by bureaucracy: the organization's main role in Jordan is the distribution a daily bread ration and a stipend of about $28 a month to each refugee family (both in the camps and living in Jordanian communities), which allows the families to buy food at local grocery stores. In the process, the organization must coordinate not only the thousands of refugees, suppliers, and grocery stores around the camp and throughout Jordan, but also work with the over 30 other relief organizations (UN related and not) seeking to help the WFP's charges in other capacities[33].

This case study will explore the design, implementation, and deployment process for the WFP's Building Blocks blockchain program. We chose this case because we believe it exemplifies a well-executed process from concept to deployment. However, it is a work in progress.

---

[28] An abridged version of this case study is available in the paper published by the Blockchain Trust Accelerator
[29] Food and Agriculture Organization of the United Nations, *SOFI 2017 - the State of Food Security and Nutrition in the World* (Paris: Food and Agriculture Organization of the United Nations,[2018]). 2.
[30] "Global Trends: Forced Displacement in 2016,"accessed Apr 23, 2018, http://www.unhcr.org/statistics/unhcrstats/5943e8a34/global-trends-forced-displacement-2016.html. 5.
[31] "Figures at a Glance," accessed Apr 23, 2018, http://www.unhcr.org/figures-at-a-glance.html.
[32] "UNHCR Jordan Factsheet: Azraq Refugee Camp (January 2018)," last modified Jan 1, accessed Apr 23, 2018, https://reliefweb.int/report/jordan/unhcr-jordan-factsheet-azraq-refugee-camp-january-2018. 1.
"Zaatari Refugee Camp - Factsheet, February 2018," last modified Feb 14, accessed Apr 23, 2018, https://reliefweb.int/report/jordan/zaatari-refugee-camp-factsheet-february-2018. 1.
[33] Ibid, 2.

As such, we will be forthright in pointing out both the strengths of the deployment process as well as the areas where the Building Blocks solution still has ground to cover before we consider it a template for others to follow.

The story

In June of 2016, one of the WFP's finance officers, Houman Haddad, approached the innovation arm of the organization and proposed that the blockchain could address a number of challenges that the WFP faced when delivering aid. The financial challenges Mr. Haddad intended to address mostly dealt with reducing interaction with and payments to banks: the goals were to limit bank transfer fees (which are typically between 1.5 and 3% of assets transferred), to limit the risk associated with advancing large sums of money to banks in countries with unstable financial systems, and to eliminate the need to share private beneficiary information with each of those banks. More promising still, Mr. Haddad proposed that the blockchain could facilitate coordination between the WFP and other relief agencies (both within the UN and otherwise), and potentially help many beneficiaries build up a verifiable financial history and basis for legal identity.

After evaluation of the technology and consultation with technical partners, the cash-based transfer (CBT) and finance divisions of the WFP determined that the blockchain offered the capabilities of a traditional solution that they cared about while providing additional strengths and optionality. In particular they thought that the blockchain provided a means to effectively coordinate the provision of aid by disparate relief agencies by facilitating reconciliation of information using a transparent record. This in turn, would reduce the incidence of redundant efforts and allow the UN to do more with existing resources.

Technical points of the solution

As the Building Blocks program was initially only used by the WFP, its primary function was at first limited to facilitating the distribution of cash transfers to the refugee population under the WFP's care. With the aid of technical partners selected through a competitive tender, Building Blocks was designed as a private fork of Ethereum for the simple reason that Ethereum appeared the most promising of the limited options available in 2016. Further, a private fork was the option that provided the chain with the greatest ability to scale as transaction volumes increased. Despite their continued success using the private fork of Ethereum, Building Blocks' creators have kept flexibility in mind from the start, and believe a switch to another chain would be relatively straightforward.

The product was designed to store no personally identifiable information (PII) on the chain: each refugee family was identified only by their UNHCR case number and verification was accomplished using a system of iris scanners that were already in place in the markets in the region. While iris scanners are being used in the Jordanian camps, any form of biometric identification or other form of authentication, such as one-time passwords, could work.

At the beginning of each month, a number of tokens are credited to each family's account. Those credits can be spent at any of the participating markets, and smart contracts are used to

transfer the appropriate number of credits from a refugee's public key to the corresponding market's public key.
At the end of the month, the WFP transfers cash directly from a WFP account to the market's account to square the balance: cash does not actually touch the blockchain. From the perspective of both the refugees and the markets, the experience is mostly unchanged. Even today, most users are totally unaware that a blockchain is involved.

Scaling

After over six months of design and implementation, the WFP began the project's deployment phase with a three-day proof of concept (PoC) in Pakistan. The goal of the PoC was to test Building Block's technical viability in the field and its ability to fulfill basic cash transfer functions. Upon satisfactory completion of the PoC, a pilot facilitating cash transfers to 10,000 refugees in the Azraq Refugee Camp in Jordan was initiated to see whether the solution could work at some scale.

Azraq was selected based on a number of considerations. First, the necessary prerequisites were in place: a local market (which is necessary if the refugees are to spend cash in the first place), and Internet infrastructure that was described as "good enough". Further, the volume of cash transferred to refugees in Jordan, and the buy in from WFP's team on the ground, meant that the pilot would help process a large amount of cash and benefit from local support.

As the Azraq pilot showed signs of success, the WFP continued to scale the Building Blocks solution: first expanding to the entirety of the Azraq camp (~30,000 people), then to the other major refugee camp in Jordan, Zaatari (over 70,000 people). As of January 2018, the Building Blocks program in the Jordanian camps is serving over 100,000 people, which represents all of the WFP's beneficiaries living in Jordanian camps.

The next step is to test whether the solution can function in a non-camp context: the WFP expects to expand Building Blocks to cover all of its charges in Jordan (over 500,000 people). In addition to scaling the project within Jordan, the WFP is currently considering additional countries in which to deploy the product later this year to prove that the solution can work in different markets. These rollouts will involve making technical changes to the blockchain, allowing users more flexibility in where they spend their cash and providing the structure to potentially include the banks themselves. The solution is currently limited to serving just over 1 million people (by Ethereum's scalability). At some point, however, Building Blocks may be used to serve many of the WFP's 80 million beneficiaries around the world.

As the deployment has grown, the WFP has become more stringent in choosing its technical partners: it has now done a full RFP and today works with two main development partners and one main advisory company (which, the WFP stressed, must be separate from the deployment partners to provide a level of objectivity); however, the organization has eight different companies under contract to provide flexibility and ease of scaling.

It is also important to note that despite the solution's success, the WFP maintained the legacy Jordanian system (based on reloadable ration cards) in reserve up until January of 2018. The

team wanted to maintain redundancy in the event that Building Blocks failed to ensure that beneficiaries wouldn't be affected.

Quantifiable outcomes

The quantifiable outcomes of Building Blocks have thus far have been remarkable: bank transfer fees have declined by 98%. The solution has also allowed the WFP to cease providing Jordanian banks with beneficiary PII and to stop transferring large sums for those banks to hold in escrow at the start of each month, thereby eliminating qualitative risks in addition to reducing costs.

While unwilling to disclose exactly how much the blockchain solution cost to implement, the organization was able to share that it has spent over $200k in the initial stages of the project and will authorize additional disbursements as they scale.

The immediate financial return has been savings of ~$40k a month in bank fees arising just from running the solution in the Jordanian camps. We were cautioned against extrapolating this figure to the rest of Jordan and internationally because savings will be context specific. That said, the organization is confident that it will save much more than $40k a month in Jordan as the solution scales, and that in scenarios deemed suitable to deploy Building Blocks, the payback period is likely to be less than a year from the perspective of CBT savings alone. The upside from the project has the potential to be even greater as the blockchain begins to facilitate coordination between the WFP and other UN agencies, thereby improving the efficiency of aid provision.

Areas for Improvement

The Building Blocks Program has drawn criticism from a few members of the blockchain community. Critics note that, as it stands, the WFP blockchain has precisely one member, the WFP itself[34]. Given that beneficiaries, markets, banks, and other UN agencies aren't a part of the blockchain network today, the project hasn't yet fulfilled all of the aspects of its original vision, and currently violates our principle that blockchains should only be used when multiple parties are submitting data. Consequently, the outcomes of the Building Blocks program thus far could, indeed, have been accomplished using a traditional database.

The WFP team is fully aware of this shortcoming and is in the process of implementing functionality that would only be possible with a blockchain, initially by on-boarding other UN agencies to facilitate coordination and reconciliation as was planned from the start: UN Women is expected to join the network in the near future, and the WFP is in discussions with a number of other UN agencies to do the same. This process will bring its own series of challenges, including answering complex questions regarding governance and the chain's consensus mechanism.

[34] David Gerard, Attack of the 50 foot blockchain, Nov 26, 2017, https://davidgerard.co.uk/blockchain/2017/11/26/the-world-food-programmes-much-publicised-blockchain-has-one-participant-i-e-its-a-database/.

Steps after that include adding non-UN humanitarian agencies to the chain to facilitate improved aid distribution, and eventually (we hope) empowering the refugees themselves by giving them more sovereignty over their data. This could help them use the UN solution as a way to build up the financial histories, credit, and verified identities that many of them lack.

Conclusion

This case study exemplifies a well-executed design, implementation, and deployment process for a blockchain concept. The initial rationale for using the technology, especially with regard to coordinating UN agencies, was sound. The design of the solution was done with the aid of outside technical partners, as we suggest, and the WFP has maintained flexibility both with regard to the platform and to it vendors.

Further, the Building Blocks pilot was implemented only after determining that pre-requisites were in place, run redundantly with the legacy solution, and was scaled in a measured fashion only upon the satisfactory completion of predefined hurdles.

Despite the impressive quantitative outcomes of the pilot thus far, as we noted, we don't believe that the solution's current incarnation illustrates blockchain's unique strengths: that will come if and when the WFP follows through on its original plan to bring other relief agencies onto the chain and then use it as a decentralized means for validation and reconciliation of data. In this regard the WFP has been limited by the small size of its blockchain team, which has led to some delays in on-boarding other UN agencies.

As one would expect of a solution intended for a large organization, Building Blocks has been forced to prove itself at each step: first in the camps, next in a country, and soon as a means for interagency cooperation. The solution is only slightly more than a year old and is currently undergoing an independent technological and strategic review; if the review goes well, we would expect adoption by more agencies will follow, and look forward to seeing how Building Blocks develops as it expands internationally and becomes a tool for collaboration. If it works, many millions of the world's most needy will be much better for it.

HEXA FOUNDATION

Swedish Lantmäteriet Land Registry Case Study

A special thanks to Jörgen Modin, Henrik Hjelte, and the rest of the ChromaWay team for their participation in and help with this case study.

What is the problem being solved?

Land is one of the bastions of the global economy not only because agriculture and its ancillary industries continue to be significant contributors to GDP (especially in the developing world), but also because land is one of the world's main stores of wealth and a vehicle to facilitate economic mobility.

According to the Bureau of Economic Analysis, privately owned land in the US alone was worth ~$21.2 trillion in 2015[35], approximately 17% more than the US' GDP in that year. The value of land relative to total GDP in less developed nations should theoretically be greater due to the lower contribution of non-land based industries (like technology). In practice, however, this is not always the case due, in part, to uncertainty around who owns the land in the first place.

As the members of Codex Labs have noted, the value of an asset, especially highly valuable illiquid assets like land or art, depends in large part on the ability of its possessor to prove ownership and provenance[36]. Further, "having title acts as verification of ownership for the current owner, which enables the ability to sell an item, insure it, borrow against it, lend it, and more"[37].

It is therefore unfortunate that the World Bank estimates that approximately 70% of the world's population "lacks access to proper land titling and demarcation"[38]. Deficient titles tend to be more prevalent in less-developed nations, but are not isolated to them: about 30% of property titles in the United States are deficient[39].

Deficient title is the source of a number of direct costs including legal and insurance fees, lengthy transaction times, and reduced liquidity. The indirect costs are even more pernicious. Unclear title places limits on economic growth and complicates the provision of aid; those unable to prove ownership of land can't use it as a source of credit and often have difficulty applying for government financial aid in the event of a natural disaster[40]. Further, deficient title creates the opportunity for land expropriation by corrupt civil servants and drives economically irrational behavior (such as underinvestment in farmland) by those who feel insecure about their

---

[35] William Larson, *New Estimates of Value of Land of the United States.* Bureau of Economic Analysis, (2016).
[36] Codex Labs, *Codex Protocol: A Decentralized Title Registry and Cryptocurrency for the Arts & Collectibles Market* (London: 2018).
[37] Ibid.
[38] Caroline Heider and April Connelly, *Why Land Administration Matters for Development.* The World Bank Independent Evaluation Group,[2016]).
https://ieg.worldbankgroup.org/blog/why-land-administration-matters-development.
[39] James Schneider Ph.D et al., *Profiles in Innovation: Blockchain, Putting Theory into Practice.* Goldman, Sachs, & Co., (2016).
[40] Gillian Tett, "Bitcoin, Blockchain and the Fight Against Poverty," *Financial Times* - 12-22, 2017.
https://www.ft.com/content/60f838ea-e514-11e7-8b99-0191e45377ec.

claim on the property. More often than not, those impacted by title issues are among the more vulnerable members of society.

This case study will explore the implementation of a blockchain land registry in Sweden by a public/private consortium led by the Swedish Lantmäteriet (which is in charge of geodata, cadastral services, and land registration). As it will show, blockchains have the potential to provide a means to secure land titles and meaningfully improve the efficiency, security, and transparency of the real estate transaction process; however, blockchains cannot solve all of the aforementioned problems alone. The implementation of a blockchain solution to address the problem involves a number of prerequisites that are typically only in place in more developed countries. As such, though it is in early stages of being deployed, we believe the Swedish implementation sets a precedent for countries with less efficient title registry systems. We hope it will provide an impetus for those countries to build out the prerequisites necessary to make blockchain solutions, land related and otherwise, viable.

The story

In 2015, the Lantmäteriet partnered with ChromaWay (a private blockchain technology company), Telia (a major Swedish mobile carrier) and Kairos Future (a consulting firm) to explore the possibility of using a blockchain to supplement the country's digital title registry system and to streamline the real estate transaction process.

The consortium's ultimate goal is to use blockchain technology to increase the transparency of the real estate transaction process and diminish inefficiency/delays by reducing redundant processes and their associated costs[41]. It is explicitly not intended to replace existing participants in the process, only to streamline it[42].

The group considered traditional alternatives to the blockchain as a means to accomplish those goals and concluded that:

"A solution with a centralised database would be less costly, but would offer vastly reduced security and minimal improvements over current processes. It is hard to compare with other solutions since we don't know of any that can provide the same capabilities and security guarantees.[43]*"*

In particular, the Lantmäteriet report cited a number of qualities of blockchains that were both attractive in the context of this particular type of implementation and difficult to replicate with traditional solutions[44]:

- Verifiably unique records, which reduce the risk of fraud
- Transparency, immutability, and resiliency, which increase security and public confidence in the system

[41] Lantmäteriet et al., *The Land Registry in the Blockchain*,[2016]), 8.
[42] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed*, (2017]), 67.
[43] Ibid, 21.
[44] Ibid, 24.

- Distributed governance, which creates clearly established and transparent processes that also improve public trust

We would add automatic reconciliation, which is not possible with separately managed traditional database instances and would meaningfully improve the efficiency of a transaction, lower the cost of processing, and reduce the risk of inconsistent records.

Technical points of the solution

The Lantmäteriet assigns each property a unique identification number linked to Sweden's digital land registry that is then stored on a private version of ChromaWay's proprietary blockchain platform[45]. For the moment, the blockchain solution is not intended to store any actual "bearer instruments[46]," and little actual data (encrypted or otherwise) is stored on the chain itself. The actual documents are stored in traditional databases by each of the entities involved in their creation, and the information on chain is mostly the digital fingerprints of those documents, or attestations by authorized individuals (like the current owner, saying the property is actually up for sale). This validated record is visible to all other entities involved in the transaction (e.g. banks)[47]. The chain's increased transparency diminishes the need for redundant verification/auditing, thereby reducing costs and expediting the transaction process. A fingerprint of the state of the private chain is periodically posted to the Bitcoin blockchain for additional security/replication[48]. In this regard, the Lantmäteriet solution is an example of an anchored hybrid chain.

As we noted in Part III, blockchains require users to have Internet access. High quality connectivity is far from a forgone conclusion in many developing nations, but Sweden benefits from excellent Internet infrastructure. Further, for title registries and applications like them, existing records must be accurate and digitized. This often requires a herculean effort and is, itself, a huge value-creation step for any nation that continues to rely on paper-based records. Sweden's property records were digitized in the 1970s[49].

A well-developed identity solution is also of critical importance in this case[50]. In the Swedish pilot, one of the nation's largest telephone companies, Telia, provides a form of digital identity. That particular solution may be suitable in this case but is unlikely to be broadly applicable[51].

As the authors of the Lantmäteriet's reports point out, Sweden needs to improve the legal and regulatory framework around digital signatures and identity before a more widely deployable identity solution can be released[52].

---

[45] Margaret Kapitany et al., *Trafi & Blockchain: An Exploration of Blockchain Data Management for the Finnish Transport Safety Agency* (Helsinki: Finnish Transport Safety Agency,[2017]), 31.
[46] Ibid, 4.
[47] Ibid, 4.
[48] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed,* 37.
[49] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed,* 17.
[50] Graglia and Mellon, *Blockchain and Property in 2018: At the End of the Beginning,* 12.
[51] Lantmäteriet et al.,*The Land Registry in the Blockchain - Testbed*; Lantmäteriet et al., *The Land Registry in the Blockchain* 22.
[52] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed*, 20.

Additionally, the consortium noted the importance of ensuring interoperability and avoiding vendor lock in.
For that reason, they didn't implement a solution that restricted to them to a particular token, and ChromaWay claims that they were selected in part because their blockchain platform is relatively easy to switch away from[53].

Quantifiable outcomes

Because the solution is still in very early stages of deployment, there is limited quantitative data regarding its effectiveness. The Lantmäteriet estimates that when fully deployed, the platform could reduce transaction times from 3-6 months to a few days[54] and save about 100mm Euro a year in the process[55]. This is not a particularly meaningful sum relative to the approximately 430bn Euro Swedish GDP, but as far as we know, the implementation is expected to be cost effective.

As the Lantmäteriet noted, the value created by a similar solution would be meaningfully greater in countries where land registries and government benefit from lower levels of trust or higher inefficiency than Sweden does. The authors of the consortium's report wrote that improved confidence in the validity of land registries would allow banks to charge lower interest rates and improve the growth of the mortgage market in developing economies. Reducing interest rates by just 0.1% "would create $14 billion per year in value" worldwide[56], and implementation of solutions like the one in Sweden may be the single "most cost efficient way to make a medium GDP per capita country into a high GDP per capita country[57]." We hope that this possibility provides yet another impetus for nations across the wealth spectrum to verify and digitize their title registries.

Scaling

After a few months of preliminary design, a proof of concept (PoC) was launched in 2016 to study the viability of blockchain technology as a means to serve the Lantmäteriet's needs.

Upon satisfactory completion of the PoC in June of that year, a pilot was initiated to implement a mobile application and to test the solution's ability to integrate with financial institutions' existing processes by executing small-scale data transfers[58]. After about two years of testing the system, the consortium is finally preparing to execute a full transaction and is in the midst of troubleshooting governance issues.

Blockchain technology's role in Swedish real estate transactions is expected to grow gradually as it shifts from being mainly an auditing mechanism and toward being the main

[53] Lantmäteriet et al., *The Land Registry in the Blockchain*, 34.
[54] Ibid, 32.
[55] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed* 17.
[56] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed* 18.
[57] Ibid, 19.
[58] Kapitany et al., *Trafi & Blockchain: An Exploration of Blockchain Data Management for the Finnish Transport Safety Agency*.

database/middleware used to store data about and facilitate the execution of transactions. This will occur as "legal, process, and security problems" are identified and addressed and as more transactions are processed using the solution in its current form[59]. The hope is that the blockchain will eventually not only facilitate "smart workflow", but also provide a means for "smart escrow" (where smart contracts can be used to facilitate transactions automatically) and potentially a full "blockchain registry" (in which the blockchain completely replaces traditional databases as a means to store actual title documents)[60].

Areas for Improvement

The consortium itself has identified three areas of weakness that are germane to our case study: the immutability of data once on the blockchain, the limited regulatory framework currently in place for blockchain-based solutions, and the problems with the current identity solution we alluded to above.

The immutability of on-chain data is problematic from the Lantmäteriet's perspective due to the regulatory requirements imposed by the EU. The consortium is still working to balance the immutability of the blockchain with the possibility of inaccurate data entry and the need to comply with European GDPR standards, including the right to be forgotten[61]. This issue is currently under consideration in EU working groups and courts.

As we have noted, when dealing with off-chain assets, blockchain network participants will often have to resort to the use of traditional recourse in the event of a disagreement. The Lantmäteriet and its partners are cooperating with other Swedish government agencies to design improved regulatory and legal frameworks to allow for the expansion of blockchain technology's role in the real estate process[62]. Those solutions are not, as of yet, in place.

From our perspective, the identity solution may be this implementation's single greatest limiting factor in that it relies exclusively on a domestic private actor and is therefore Sweden-specific and centralized; perhaps a similar partnership would function in countries where the populace has a significant degree of trust in a single institution, but those cases are limited.
The consortium's report was honest in its evaluation of this flaw, noting that "making an ID-solution and securing the adoption of the ID solution may be the most difficult problem for the solution to be used in a widespread way globally.[63]"

Conclusion

This case study has explored the design and deployment of a blockchain to supplement the real estate transaction process in a developed nation. While still in very early stages of the deployment process, the Lantmäteriet solution has thus far aligned with our principles for deployment and seems to hold promise.

---

[59] Lantmäteriet et al., *The Land Registry in the Blockchain - Testbed* 3.
[60] J. Michael Graglia and Christopher Mellon, *Blockchain and Property in 2018: At the End of the Beginning* (Washington D.C.: [2018]). 20.
[61] Ibid, 20.
[62] Ibid, 71.
[63] Ibid, 66.

Land title registry applications in general check the boxes of our decision tree, and the Swedish Lantmäteriet's design is consistent with what we would have expected from a solution involving real estate. Successful title-related implementations at this stage are likely to be supplementary to existing processes rather revolutionary. This is partly due to the technology's immaturity and partly to the reticence of existing institutions to see their influence reduced. As the technology develops further, greater decentralization and the elimination of certain parties currently involved in transaction processes (title insurers or notaries, for example), may be possible. For the moment, however, we expect that most blockchain solutions involving real estate will operate within the bounds of the existing system and often with the support/initiative of a government entity. The exception will arise in cases where trust in existing institutions is irretrievably low.

The consortium's solution is also consistent with our expectations from a technical perspective. Most blockchains used in title registry cases are likely to be a private or hybrid chain so as to allow for identification of the parties involved; identification is important for a number of reasons, not least the collection of taxes and compliance with existing KYC (know your customer) regulations[64].

We re-emphasize that the deployment of blockchain technology in this case was dependent on a number of pre-requisites, including the existence of an accurate and digitized title registry. Countries lacking those prerequisites should invest in them first. Doing so will, in and of itself, generate enormous value. The Swedish solution was deployed only with those prerequisites in place and was also aligned with our recommendation to maintain interoperability and is structured to avoid vendor lock-in.

Though government entities are likely to be involved in the majority of title-related implementations, we note that despite being a large government entity, the Lantmäteriet did not manage all aspects of the design, implementation, deployment, and maintenance of the blockchain from scratch. We expect most blockchains will follow this model and be built by some type of partnership, rather than by a single organization attempting to tackle the whole thing itself.

Lastly, while the Swedish deployment has been very conservative, it was in line with our recommendation that organizations design slowly and use caution as they scale, starting with a PoC and then proceeding to a pilot to attain quantitative and qualitative data. If the feedback is favorable, gradual scaling of the pilot by either expanding its geographic scope in its existing role or increasing the blockchains' role in transactions, may be merited.

It is our hope that through the replication of this data-driven and measured approach, blockchain's use in land title applications will become widespread: the value resulting from the use of blockchain to reduce the incidence of deficient titles could conservatively be in the hundreds of billions of dollars.

[64] Graglia and Mellon, *Blockchain and Property in 2018: At the End of the Beginning,* 14.

Summary of Case Studies

As these case studies have shown, blockchains and distributed ledger technology have already been deployed to varying degrees by a variety of organizations across the public, private, and non-profit sectors and are adding value in different ways.

Despite being of different sizes and operating in different sectors, the processes the organizations featured in this section engaged in to design and deploy their solutions were not meaningfully different. Though none of the deployments are perfect, they all contain elements of the concepts and principles we expounded on in Part III and IV and serve as useful examples both of what blockchains and DLT can accomplish.

We look forward to watching these and many of the other blockchain deployments considered for inclusion in this paper as they evolve and fulfill their potential. None embody the promise of a crypto-utopia, but if they deliver on their goals, they will all yield a meaningful improvement in the lives of their beneficiaries.

Part VI: Closing Thoughts

In this paper we have explained the technical foundations of blockchains, highlighted their strengths and weaknesses relative to their traditional database peers, and tried to provide structure to the decision-making process as organizations consider whether and how to design and deploy a blockchain solution.

We believe that blockchain technology and distributed ledgers have enormous potential, but that the technology is still early in its evolution. It is important to understand that blockchains, like all technologies, have limitations, and that even at maturity, they may be unsuited for certain applications. In some cases a traditional database solution may prove cheaper and easier to implement and yield better results; however, there are many applications where traditional databases have proven ineffective, inefficient, or unsecure. In these cases, blockchains may be able to add considerable value. Thus, experimentation with the technology, despite its relative immaturity, can yield considerable benefits, as illustrated by early adopters of the internet during the 1990s.

Further, we believe that when deployed in the appropriate scenario, blockchains' unique technical qualities will help create extraordinary value in the public, private, and non-profit sectors. Blockchains don't have to dominate the global economy in order for the world to benefit from the technology: they can create considerable value in supplementary roles and by addressing problems where traditional solutions have proven incapable. Whether they begin to replace traditional solutions to a greater degree depends on a number of factors, among them the outcome of the investment and research being pursued by members of the blockchain community.

Closing thoughts

The human race has succeeded due to our ability to collaborate. Unfortunately, cases arise when we can't find partners to collaborate with, or where our partners may choose to achieve gains at our expense. The result is that we are prevented from working toward ends that may benefit us all. Blockchains and their DLT peers provide one means to solve some of these protracted problems. By facilitating trust and coordination, the technology can help create opportunities for collaboration that were previously rendered impossible by information asymmetry or digital and physical boundaries.

The early beneficiaries of the technology are likely to be established corporations who are able to use blockchains and other forms of distributed ledger to improve profitability and the experience of their customers. However, in the long run, we hope that the biggest beneficiaries of blockchain technology will be individuals rather than companies; DLT is, after all, a theoretically open and democratic technology. Thus, the benefits of blockchains' growth, both in the form of reduced costs and the opportunity for new economic activity, should be widespread.

With those benefits will likely come some disruption: if blockchains gain traction, it is likely that some jobs will be displaced and that some intermediaries will see their roles shrink. Despite this dynamic, the technology's benefits from both an economic value and social impact

perspective have the potential to be significant. With the appropriate infrastructure, blockchain offers the possibility to facilitate the inclusion of billions of people who have been left behind because including them using legacy systems is impracticable or economically untenable. Blockchains may also allow for greater coordination within and among organizations across sectors to reduce redundant and ineffective spending.

From a governance standpoint, greater transparency and higher accountability will empower individuals to make better-informed and more independent decisions about how they manage their data, as well as provide improved insight into the inner workings of the institutions that people trust. The efficiency gains to digitizing records and managing government services through interoperable sets of data could rekindle citizen's faith in the ability for government to solve 21st century problems. As such, though they could theoretically reduce our reliance on centralized authority in certain cases, blockchains may also provide a means for institutions to prove their worth and engender new trust.

All that said, blockchains aren't a panacea. As we have noted, the technology has inherent limitations that we believe will keep certain aspects of human interaction forever beyond the reach of code. This is as much a hope as the outcome of our research: a little irrationality keeps life interesting.

HEXA
FOUNDATION

Acknowledgements

This paper was only possible due to the patience and kindness of the many people who shared their time and knowledge with us in person, over the phone, and by email. If the research herein amounts to anything, it is because of them and the openness of the blockchain community. Thank you all:

Abhay Gupta, Ad Kroft, Adithya Kumar, Alex Tong, Andrew Bakst, Ankur Kapadia, Ari Juels, Aymard Dudok de Wit, Baloko Makala, Ben Burke, Ben Joakim, Bob Visnov, Brian Waterhouse, Charlie Knoll, Chris Wood, Christopher Verceles, Cole Riccardi, Cooper Schorr, Daniel Yim, David Huysman, Drew Sadler, Eliot Hedman, Emily Mitchell, Ethan Schmertzler, Felipe Daguila, Fennie Wang, Frank Yiannas, Ish Goel, Jaka Jaksic, Jamie Green, Jess Houlgrave, Jörgen Modin, Joseph Bonneau, Ken Marke, Kevin Hu, Megan Cojocar, Mert Ozdag, Mike Kalomeni, Nick Martitsch, Nick Miller, Pablo Tutino, Philip Gradwell, Priyanka Dekhar, Pulkhit Agarwal, Rachel Pipan, Raphael Mazet, Richard Li, Riley Hughes, Russell Yanofsky, Sam Ng, Shaillee Adinolfi, Stephanie Seale, Stephany Zoo, Stephen Hadeed Jr., Sterling Johnson, Tejas Bhatt, Titus Capilnean, Zhuzeng Lu.

Special thanks are owed to the following:

- Titus Capilnean, Felipe Daguila, Priyanka Dekhar, Philip Gradwell, Kevin Hu, Jaka Jaksic, Zhuzeng Lu, Alex Tong, Bob Visnov, and Daniel Yim, for their input and time
- Netta Korin for her willingness to lend a helping hand to anyone in need and making sure this paper saw the light of day
- Ran Melamed for his excellent counsell and hours spent editing
- Ethan Chernofsky, for his work getting this paper into the hands of readers
- Tomicah Tillemann and Allison Price for their input and advice
- Tejas Bhatt, David Davies, Jamie Green, Nick Miller, Jörgen Modin, Sam Ng, Drew Sadler, Frank Yiannas, and the rest of the teams at AgUnity, Walmart, and the World Food Programme, for their patience answering hours of questions and countless emails

Best,

Alex

HEXA FOUNDATION

Glossary

Note: The taxonomy of the blockchain ecosystem is still evolving. While we believe these definitions to be accurate, they may not be the ones used by other members of the blockchain community and encourage our readers to ask for clarification when someone uses a technical term.

Anchoring: An auditing mechanism used by some private chains that involves periodically storing a cryptographic fingerprint (called a hash) of all of the data on a private chain on a public platform (like Bitcoin or Ethereum) as a means to increase security.

Append-Only: A quality of a database such that users can only write to, but not delete from the data.

Bitcoin: The first blockchain technology, designed by Satoshi Nakamoto and launched in 2009.

bitcoin (BTC): The cryptocurrency used on the Bitcoin blockchain.

Block Reward: A prescribed number of cryptocurrency tokens awarded to the node that successfully mines a block in a blockchain. These provide economic incentive for nodes to participate in a blockchain and form part of the consensus mechanism.

Cryptocurrency: A type of digital currency, created using cryptographic techniques, which is used within a particular blockchain ecosystem.

Decentralized Application ("DApp"): A smart contract or series of smart contracts with an unbound number of participants that may or may not involve digital assets and is used to accomplish a goal (like sharing files) on a peer-to-peer network[65].

Distributed Ledger Technology (DLT): A means of data amalgamation and distribution that allows data to be synchronized and shared across a peer-to-peer network[66]. Blockchains are a type of DLT, but not all DLT is a blockchain.

Double Spending: Attempting to spend the same resource (like a bitcoin) twice.

Ether: The main cryptocurrency used on the Ethereum blockchain.

Ethereum: A public blockchain platform designed by the Ethereum Foundation and launched in 2015.

---

[65] Ethereum blog, May 6, 2014, https://blog.ethereum.org/2014/05/06/daos-dacs-das-and-more-an-incomplete-terminology-guide/.

[66] Harish Natarajan, Solvej Karla Krause and Helen Luskin Gradstein, *Distributed Ledger Technology (DLT) and Blockchain* (Washington DC: The World Bank,[2017]). http://documents.worldbank.org/curated/en/177911513714062215/Distributed-Ledger-Technology-DLT-and-blockchain.

Fault Tolerance: The ability of a database to continue to function despite the failure of multiple nodes in its network.

Fork: A split in a blockchain typically involving a change to the underlying protocol. Following a fork, network members choose which set of protocols to follow.

Friction: "Costs, both implied and direct, associated with a transaction"[67].

Full Node: A node that stores a full copy of the blockchain database and, typically, acts as a validator.

Hash: A unique and standard-length string of letters and numbers that can be derived from any piece of data using a pre-existing cryptographic algorithm.

Hybrid Blockchain: A type of blockchain that combines elements of private and public chains. These typically take one of two forms: a private blockchain built on a public platform like Ethereum (sometimes called a "side chain"), or a private chain that is "anchored" to a public platform.

Latency: The amount of time it takes data to propagate across and be stored by a network (used to describe the speed of a particular network).

Lite Node: A node that stores only the portion of the blockchain database that is relevant to its owner.

Metadata: Data that describes other data and is used to organize and locate data within a system (e.g. the time that a message was sent).

Miner: A full node that is invested in a blockchain and attempts to propose the next block in said chain as prescribed by its protocol.

Network: An interconnected system of two or more digital devices that can exchange data.

Node: Another term for a computer that is a member of a distributed ledger or other type of network.

Off-chain: Data that originates from an event or process that takes place outside of the confines of the blockchain and is not governed by its protocol.

Openness: The degree of ease or difficulty for a node to join a network.

Orphaned Block ("Orphan"): A block of data, which may or may not be valid, that is not included in the blockchain's prefix.

[67] "Definition of "Friction Costs" - NASDAQ Financial Glossary," "Definition of "Friction Costs" - NASDAQ Financial Glossary," accessed Apr 17, 2018, https:/www.nasdaq.com/investing/glossary/f/friction-costs.

Peer-to-Peer Network: A network of that allows participating nodes to exchange information with each other without relying on a central node as a relay.

Prefix: Otherwise known as the “longest chain”, the prefix is considered the consensus view of a blockchain network at any one time.

Private Blockchain: A blockchain in which a single entity, or group of entities, that is typically already part of the network, select what nodes have the ability to read, write, and validate data added to the chain. Data is typically visible only to those allowed into the network. Sometimes also called a “permissioned” blockchain.

Private Key: A long string of randomly generated alphanumeric characters that is cryptographically linked to a public key and functions as a password to generate a signature that can be used to authorize transactions and authenticate data.

Proof of Stake: A consensus mechanism used in blockchains that is based on a miner’s existing token/cryptocurrency holdings.

Proof of Work: A consensus mechanism used in blockchains that is based on the computing power that a miner contributes to a particular blockchain.

Protocol: 1) Another name for a particular blockchain or distributed ledger; 2) A difficult to change rule or set of rules that defines the terms by which a blockchain or other type of DLT operates.

Public Blockchain: A blockchain in which read, write, and validate permissions are theoretically open to anyone with access to the Internet and the appropriate hardware. Data is typically visible to anyone who joins the network. Sometimes also called a “permissionless” blockchain.

Public Key: a long string of randomly generated alphanumeric characters that is cryptographically linked to a private key and functions as a form of address or alias. This is also known as an “address.”

Public Key Cryptography: A cryptographic method used to securely exchange information without revealing the information itself.

Scalability: The capability of a system, organization or process to sustain or increase its performance and accommodate growth.

Sharding: A technical term used to describe the partitioning of a database into smaller parts.

Side Chain: A type of hybrid blockchain in which a private blockchain is attached to a public blockchain like Ethereum.

Signature: A cryptographic demarcation indicating that a particular piece of data is validated by a certain user using their private key without revealing the private key itself.

Smart Contract: A program written onto a blockchain that performs a specific task once predetermined conditions are met. Once rules and penalties are agreed to by its parties, a smart contract becomes a self-executing and self-enforcing contract.

Tamper Resistance: Used to describe the quality of blockchains that makes data encoded in them hard to change and makes changes, when executed, evident.

Token: A type of cryptographically secured asset used for different applications on blockchains (it may or may not have monetary value). These can either be used as an endemic currency within the platform, or represent assets in the real world such as electricity, financial credit, or physical space in a shipping container. Also referred to as a "coin".

Tokenization: The conversion of the rights to a real world asset (digital or physical) into a digital token on a blockchain that can be bought and sold[68].

Transaction: An exchange of value or data in a blockchain network. That data can be the actual, unencrypted information, its encrypted version, or a cryptographic fingerprint that represents the data while the data itself is held off-chain.

Transaction Fees: The fees that network members pay to incentivize miners to verify transactions and include said data or transaction in their block.

Valid: A block or transaction that is aligned with the rules established by a particular protocol.

Validator: A node that stores a full copy of the blockchain ledger and is authorized to validate data in a blockchain network.

Note: Some of the definitions in this glossary were first published in a paper Alex Knight co-authored with Tomicah Tillemann, Allison Price, and Gloriana Tillemann-Dick. which can be found at: https://www.newamerica.org/bretton-woods-ii/blockchain-trust-accelerator/reports/blueprint-blockchain-and-social-innovation/. These definitions are being reproduced here under the Creative Commons Attribution 4.0 International License (https://creativecommons.org/licenses/by/4.0/legalcode)

[68] "How Tokenization is Putting Real-World Assets on Blockchains," last modified March 30, accessed May 10, 2018, https://www.nasdaq.com/article/how-tokenization-is-putting-real-world-assets-on-blockchains-cm767952.