OPEN FOR FEEDBACK

How can smart contracts be useful for businesses?

Open Data Institute

About

This report has been researched and produced by the Open Data Institute, and published in April 2018. Its lead authors were Jamie Fawcett, Jared Robert Keller and James Maddison, with contributions from Caley Dewhurst, Anna Scott, Olivier Thereaux, Peter Wells and Jeni Tennison. The report was produced in collaboration with Navin Ramachandran, UCL Centre for Blockchain Technologies and IOTA Foundation, and James Brogan, MD candidate at Albert Einstein College of Medicine and Research Fellow at UCL Centre for Blockchain Technologies.

If you would like to send us feedback or comment on this document, please get in touch by email at RandD@theodi.org.

Contents

Executive summary

Smart contracts and distributed ledgers might be useful for enabling efficient, trusted interactions, but it all depends on context and design.

All daily interactions between people, businesses and other organisations are underpinned by trust – each party trusts that the other will behave in a certain way. Trust is therefore central to the functioning of society and the economy. Technology has long played a part in informing trust in interactions, particularly in recent decades. The latest set of emerging technologies that have been mooted as having the potential to play a role are distributed ledgers and, in particular, smart contracts. In this report, we examine the potential of distributed ledger technologies to inform trust, focusing on the role of smart contracts.

Part 1 examines the potential of distributed ledgers and smart contracts to underpin trust within interactions. Our goal in Part 1 is to provide crucial information about these technologies and raise important questions about their benefits and limitations, in order to help businesses ask some of the right questions when attempting to decide whether distributed ledgers and smart contracts can be deployed within their business to solve real-world problems.

Part 2 explores the various ways in which distributed ledgers and smart contracts can be deployed depending on the use case, industry context and the needs of the various parties involved. Our goal in Part 2 is to help businesses that are keen to pursue smart contracts define the type of system that will prove most beneficial for their business and their clients.

Part 1: Are smart contracts useful for me and my business?

The original blockchain was designed to enable financial transactions without the need for any trusted third party. The design of the system relies on storing records of all transactions on a new kind of database, with a unique set of properties, that engender trust between members of the network in those transactions. In particular, this trust resides in the distributed nature of the database – with every member of the peer network having a copy of the blockchain and equal authority to add to it. With no central copy, every member of the network, or node, can add to the database, though they must reach consensus before doing so – a process typically handled through the use of cryptography and economic incentive. Blockchain databases can be public, so anyone can join, or can have restricted permissions to read or write to the database, depending on the design chosen.

Distributed ledgers such as these can arguably underpin trusted exchange of cryptocurrencies or other financial assets in the absence of trusted third parties rather effectively. However, many people are excited by their potential application to other use cases – typically using the immutable and distributed nature of the database to create a verifiable, single, trusted record of particular events. This can open up the possibility of answering some of the challenges that cannot be solved with a centralised database – primarily because businesses would no longer need to trust a single third party to operate the system or database. Examples of how distributed ledgers could be applied include creating a single register for art and collectibles (Codex protocol), guaranteeing the integrity of digital archives (ARCHANGEL) or managing digital music rights (Blokur).

When asked to identify the most promising use cases for distributed ledgers and smart contracts, many of the people we spoke to immediately reached for the

example of supply chains – primarily looking at distributed ledgers and smart contracts as a way to solve insidious supply chain issues such as sustainability, counterfeiting and child labour. The idea is that creating a single, transparent, authoritative, immutable ledger of goods moving through a supply chain, which all members of that supply chain can trust, might help to track the provenance of goods. Examples of companies attempting to tackle these issues using distributed ledgers include Everledger, who aim to certify the provenance of high-value assets such as diamonds, and arc-net, who aim to trace goods from source through production to consumer.

Smart contacts offer an extension of distributed ledgers, providing the means to automate processes within the database while retaining characteristics of distributed ledger technologies, especially around immutability and the creation of trust. A smart contract is a piece of executable computer code stored on a distributed ledger that, when certain conditions are met, can automatically modify data on that ledger. Smart contracts potentially offer a way to enable transparent, auditable and efficient interactions between people, businesses and governments – especially in complex networks or industries with many different players where using a trusted central authority or marketplace to approve, administer and record the interactions is difficult.

Again, the use case for supply chains can help detail the potential of smart contracts, especially when it comes to automating processes around things like certification and authenticity. Examples of companies attempting to tackle these issues include Provenance, who are using smart contracts to automate verification of certification, and Sweetbridge, who aim to use smart contracts to automatically trigger work orders.

There are places where smart contracts and distributed ledgers might be useful in tackling business problems, but businesses need to identify a clear use case where there is lack of trust between multiple actors and no central authority is trusted to administer the entire system.

Part 2: what type of smart contract system should I pursue?

Once a problem has been clearly defined, and there is the possibility that smart contracts and distributed ledgers might be useful, then businesses will need to design an approach which takes into account not only the features and limitations of the technology, but the features and limitations of the industry or system in which the problem exists. We identified three key challenges that businesses need to consider and decide how to tackle. Businesses will need to:

i) Identify how to ensure that data gets into the ledger in a trustworthy manner

Almost all use cases for smart contracts rely on external data in order to execute their terms. That external data must be written to the ledger in a way that all members of the network can trust, otherwise the smart contracts are unlikely to be useful.

ii) Identify how to deal with edge cases and resolve disputes

Regardless of how well smart contracts are implemented, there will always be errors and bugs in code. This means there will likely be edge cases where smart contract execution will result in disputes, which will need to be resolved to the satisfaction of most members of the network if they are to trust the system.

iii) Identify how the system will be funded and governed, and how value will be transferred

> All technology systems and business processes create setup and operational costs. How these costs are administered and by whom will have a big impact on not only trust in the system but incentive to participate.

Businesses need to make decisions about the extent to which they can use the technology and the extent to which they can use existing industry mechanisms to tackle these three challenges. In this choice between idealism and pragmatism is the implicit challenge of designing solutions that fit the needs of the business and the industry without losing all the potential benefits of this technological approach. Such a balance is difficult, and businesses need to be mindful about the different approaches that can be taken. Given the reliance on industry context, many systems will require different approaches in response to different challenges, and often these will be a mixture of the ideal and pragmatic.

Key takeaways

Daily interactions between people, businesses, and other organisations are all underpinned by trust.

Distributed ledgers are an emerging set of database technologies that have the potential to play a part in informing this trust – using their unique properties of immutability and distributed maintenance to create a verifiable, single, trusted record of particular events.

Smart contracts are pieces of executable computer code stored on a distributed ledger that, when certain conditions are met, can automatically modify data on that ledger – potentially providing the means to automate different processes within the database.

Distributed ledgers and smart contracts are potentially useful for businesses but only if there is a clearly defined use case where there is lack of trust between multiple actors and no central authority is trusted to administer the entire system.

Having identified a clear use case, businesses looking to implement distributed ledger and smart contract approaches to tackle challenges need to remain pragmatic about the capabilities of the technology and the existing industry context.

Specifically, businesses should be prepared to make decisions about:

> **i) how to get data into the ledger in a trustworthy manner**
> **ii) how to handle edge cases and resolve disputes**
> **iii) how to fund, govern and administer the system effectively**

About this report

In 2016, the Open Data Institute (ODI) carried out research into the application of blockchain and distributed ledger technologies, publishing the results in 'Applying blockchain technology in global data infrastructure'. In that report, we examined the basics of blockchains and distributed ledger technologies, potential use cases and applications beyond FinTech, and some of the challenges facing those attempting to tackle problems using these technologies. Our goal was to help business leaders and policymakers make informed decisions around their adoption of these new technologies.

Since we released that report, blockchain and distributed ledger technologies have continued to develop at a rapid rate. More tools and technologies have become available, and more companies and organisations have begun to develop products and services. As these technologies appear to grow more mature, interest from businesses about applying them to real-world problems grows too. The main challenge for businesses is identifying whether the problem they are trying to tackle requires, or would benefit from, the application of distributed ledger technologies. In this report we move beyond the potential challenges raised in our first report by focusing specifically on smart contracts and the role they might play in informing trust between people and organisations.

As with our previous report, we have primarily focused on the potential of distributed ledgers outside the finance sector, because the financial services industry is already heavily invested in potential applications in their industry. Our goal is to provide an overview of the potential for businesses in other sectors, and for cases that go beyond support for cryptocurrencies. To support this approach, we have used a single prominent use – supply chains – as a frame of reference to highlight not only potential applications of distributed ledgers and smart contracts but also three of the main challenges businesses will face when attempting to implement such systems.

Even with recent developments, it is perhaps still too early to pass judgement on individual technologies and applications. In this report, as with our last one, we instead raise the questions about which we think companies should be aware. It is our hope that by raising these questions we can provide businesses with a critical lens to evaluate the value of these technologies in their specific use cases.

Part 1: Smart contracts and uses for business

Interactions between people, businesses and other organisations are underpinned by trust. This report examines what role, if any, smart contracts and distributed ledgers can play in informing this trust.

Every day, billions of people, businesses and other organisations interact with one another – they make exchanges, enter into agreements and come to understandings. All these interactions are underpinned by trust – each party expects and trusts that the other will behave in a certain expected way, following a set of implicit or explicit rules. Trust is central to the functioning of society and the economy; without it, businesses, governments and people lack the confidence to meaningfully interact.

In modern societies and economies, trust is informed by a wide range of different factors and derived from a variety of different sources. The rule of law, protections of regulators, economic incentives, reputation and brand recognition, and many others, all contribute to the trust we place in certain interactions. The sources we rely on, and the extent to which we rely on them, depend on the type of interaction. Relationships between large multinationals, for example, are more likely to be governed by trust in strict legal arrangements, whereas a consumer's choice of a service provider is more likely underpinned by trust in the reputation of the brand.

Technology has been one of these sources underpinning trust for a significant amount of time. For example, the creation of double-entry accounting and the printing press helped engender trust in trade in 14th-century Europe.[1] Since the late 20th century, digital and data technologies have been playing an increasingly large role in underpinning trust in transactions. For example, secure web browsing and antivirus software give us confidence while shopping or banking online. The goal of many of these technological innovations is to enable trusted transactions to occur more efficiently, often by making them faster and easier than existing processes. Lately, one set of emerging technologies that have been touted as having a potential role in informing trust are blockchains and distributed ledgers.

Blockchains and distributed ledgers are emerging technologies that can be used to store, and in some cases manage, data. Data is essential for the modern age; it is infrastructure for the whole economy – and blockchain and distributed ledger technologies form part of our data infrastructure.[2] Data infrastructure consists of data assets, the organisations that operate and maintain them, and guides describing how to use and manage the data. It includes people, processes and technology.[3] As technologies that store data, it is important to understand what role blockchains and distributed ledgers might play in creating a robust data infrastructure.

---

[1] Geofrey T Mills (1994), 'Early Accounting in Northern Italy: The Role of Commercial Development and the Printing Press in the Expansion of Double-Entry From Genoa, Florence, and Venice', http://www.accountingin.com/accounting-historians-journal/volume-21-number-1/early-accounting-in-northern-italy-the-role-of-commercial-development-and-the-printing-press-in-the-expansion-of-double-entry-from-genoa-florence-and-venice

[2] James Smith, Jeni Tennison, Peter Wells, Jamie Fawcett, Stuart Harrison (2016), 'Applying blockchain technology in global data infrastructure', https://theodi.org/article/applying-blockchain-technology-in-global-data-infrastructure

[3] For more information, see: https://theodi.org/topic/data-infrastructure

In this report, we examine the potential application of blockchain and distributed ledger technologies to inform trust. In particular, we focus on the role of smart contracts – programmable and executable computer code stored on distributed ledgers that offer a means of codifying and automating complex interactions. Our aim is to offer neutral, unbiased advice to businesses to help them understand smart contracts and some of their potential applications. Some businesses may be exploring whether to begin building smart contract systems, whereas others may be deciding whether to join an existing system. This report is aimed at both audiences and we hope it will enable businesses to make informed decisions about using or investing in smart contracts as a solution to real-world problems.

Blockchains, distributed ledgers and trust

Blockchain emerged in response to a specific type of interaction use case – financial transactions. In traditional financial transactions, the exchange of value between any two parties, be they people or businesses, relies on one or more trusted third parties. The implicit third party in most transactions is a government or central bank that issues the currency used to exchange value. In most transactions, especially those between businesses, the parties will also rely on one or more other organisations to enact the transfer, for example the bank or banks with which they hold accounts.

The original blockchain, the Bitcoin blockchain, was designed to enable financial transactions without the need for any trusted third party. Driven in large part by an ideological rejection of value exchange underpinned by state institutions, the idea was to provide a medium for parties to exchange and transact directly with one another. The design of the system relied on storing records of all Bitcoin transactions on a new kind of database, with a unique set of properties, that engendered trust in those transactions.

In our previous report, we outlined the unique set of properties of these new databases and explained how these properties arise from the design of the database. Below, we have provided a brief overview of the key properties of distributed ledgers that enable them to inform trust in transactions.

Key properties of distributed ledgers

These properties are drawn from Greenspan (2015) ‘Avoiding the pointless blockchain project’.[4]

- **Shared read:** Blockchains are a structured data store that many people can read.
- **Shared write:** As well as read, many people can write data into the database.
- **Absence of trust:** The different writers do not have to trust each other not to manipulate the shared database state.
- **Disintermediation:** There is no need for a trusted intermediary to enforce access control.
- **Transaction interaction:** Records in the database depend on, and link to, each other.
- **Validation rules:** Rules around database transactions are well defined, such that anyone with a copy of the database can validate that it has been maintained correctly.

[4] Gideon Greenspan (2015), ‘Avoiding the pointless blockchain project’, https://www.multichain.com/blog/2015/11/avoiding-pointless-blockchain-project/

Important characteristics of distributed ledgers

Distributed ledgers have no central storage location, no 'primary' copy, but rather are maintained by a peer network of nodes; every node has a copy of the blockchain and has equal authority to add to it. In order to ensure that only one node may add a block to the end of the chain at any given time, blockchains need a resolution mechanism to decide which block is submitted to, and accepted by, the network.

There are a number of different resolution mechanisms or 'consensus algorithms' that can be built into distributed ledgers. Choosing a consensus algorithm for a distributed ledger can greatly affect the scalability of the system – in particular determining the amount of energy and computing power required to support its operation. The current most common consensus algorithm is 'Proof of Work', used by both the Bitcoin blockchain and the Ethereum blockchain, although the well-documented high-energy usage has led many to explore alternative mechanisms.[5] We have outlined several different types of consensus algorithm below.

Consensus algorithms[6]

Each consensus algorithm determines which node gets to add its data to the database. The data being added is a combination of all the transactions submitted to the network over a given period.

- **Proof of work:** nodes must compete to prove they have solved a complex cryptographic puzzle and are rewarded for doing so – a process referred to as 'mining'
- **Proof of stake:** nodes must prove they own a certain amount of cryptocurrency; the more they own the more likely they are to be chosen to provide the next block – a process referred to as 'validation'
- **Proof of activity:** nodes must first undergo a proof of work cryptography solution, which is then 'validated' through a proof of stake mechanism

These consensus algorithms often offer incentives for nodes to participate in the process of creating a new block, often in the form of cryptocurrency. They rely on this incentive to encourage many nodes to participate in the process in order to help guarantee the integrity of the system. This approach is taken for many distributed ledgers that are public – anyone can view, participate in a transaction or run a node, such as the Bitcoin blockchain or the public Ethereum blockchain.

However, recently there has been a trend towards creating networks that require permission to join. While they still operate consensus algorithms, the 'rewards' do not need to have economic value, assuming that those participating will choose to run a node. While these systems create permissions, and in some cases access may be determined by a single actor, there are still potential benefits in distributing the running of the system for creating trust between participants in that system.

---

[5] Christopher Malmo (2017), 'One Bitcoin Transaction Now Uses as Much Energy as Your House in a Week', https://motherboard.vice.com/en_us/article/ywbbpm/bitcoin-mining-electricity-consumption-ethereum-energy-climate-chang

[6] Amy Castor (2017), 'A (Short) Guide to Blockchain Consensus Protocols', https://www.coindesk.com/short-guide-blockchain-consensus-protocols/, See also Arati Baliga (2017), 'Understanding Blockchain Consensus Models', https://www.persistent.com/wp-content/uploads/2017/04/WP-Understanding-Blockchain-Consensus-Models.pdf

Permissions[7]

Distributed ledgers can be classified by who has access to read and write to the database:

- **Public:** anyone may have a copy of the database and anyone may write to it, sometimes referred to as 'permissionless'
- **Permissioned:** Anyone may have a copy of the database, but only certain parties may write to it
- **Private:** Only certain authorised users have access to the database, whether for reading or writing

Trust beyond transactions

When the first blockchain appeared in 2009, the ability of distributed ledgers to underpin cryptocurrency transactions in the absence of trusted third parties appealed to a variety of different people and organisations, in particular the finance industry. Soon after, other organisations and industries began to explore how the unique properties of this new technology could play a role beyond its immediately intended use case. In particular, there was interest in using the immutable and distributed nature of the database to create a verifiable, single, trusted record of particular events. The hope was that by attaching various other types of data to a particular transaction - eg time, location, quantity, ownership history - organisations could use the properties of distributed ledgers to create verifiable, timestamped records.

Using distributed ledgers in this way made it possible to dream of applying the technology to a wide range of use cases beyond cryptocurrencies. From supporting the verification of the content of digital files to recording the ownership or transfer of assets that are not stored on the ledger, many have argued that distributed ledgers are capable of playing a role in underpinning trust in a variety of different use cases.

Below is a small sample of the companies and projects using distributed ledgers to create verifiable records that underpin trust within their respective industry.

Applications of distributed ledgers beyond cryptocurrencies

Codex protocol[8]

Unlike most other asset classes, the art and collectibles market does not have a central title registry. This makes determining ownership of assets difficult. Records of past ownership are important to collectors as they help to prove the authenticity of an asset and its accurate valuation. Codex Labs have created the Codex Protocol Title Registry - an immutable, decentralised title registry that aims to engender trust around the origins of art and collectibles, and in turn create a fairer market for collectors.

ARCHANGEL[9]

Archived documents, images and videos enable future generations to understand society as it exists today. However, as organisations continue to shift towards digital practices, discerning whether these files have been altered will become increasingly difficult. The ARCHANGEL project is exploring the use of distributed ledger technologies as a mechanism to verify that digitally archived files have not been modified. By storing cryptographic

[7] For more details see: https://theodi.org/article/applying-blockchain-technology-in-global-data-infrastructure/
[8] For more information about Codex Protocol, see: https://www.codexprotocol.com/
[9] For more information about ARCHANGEL, see: http://blockchain.surrey.ac.uk/projects/archangel.html

hashes of the content of each file, the ARCHANGEL project aims to allow organisations to guarantee the integrity of archived content over time, and allow people to authenticate that content in an archive has not been changed beyond necessary format shifting.

Blokur[10]
Music publishing can be complex. With songwriters, recording artists and record labels all playing a role in the production and publishing of songs, determining who the legal owner of a particular composition is can be difficult, even for the parties involved. Being able to determine legal ownership is important, as individuals and organisations receive royalty payments for both sales of an original composition and sales of any cover recordings based on that original song. Blokur are using distributed ledger technologies to create an immutable record for the ownership of music rights by sourcing data from music publishers and the collective management organisations that represent artists. Using machine learning algorithms, Blokur identifies and resolves discrepancies about ownership rights in this data, eliminating the need to involve expensive third parties.

The supply chain use case

In the last report, we determined that much of the excitement around distributed ledgers was focused on the technical capabilities of the technology and not on the problems that the technology might be used to solve. With this report we want to understand the promise of distributed ledgers and smart contracts in the context of a specific use case in order to interrogate their usefulness in practice.

When asked to identify the most promising use cases for distributed ledgers and smart contracts, many of the people we spoke to in this space immediately reached for the example of managing supply chains. In particular, they focused on how smart contracts and distributed ledgers may play a role in connecting disparate members of a supply chain and helping to trace the movement of goods as they pass between members of that supply chain.

Given the widespread interest in this use case, and the number of companies and initiatives trying to address this challenge, we chose to use it as a lens through which to examine the potential benefits and limitations of applying distributed ledgers and smart contracts to tackle real-world problems.

Yet while the relevance of distributed ledgers and smart contracts to the supply chain use case is broadly agreed upon, companies attempting to implement distributed ledgers and smart contracts are taking a range of different approaches. This is in large part because the technology is still emerging and the platforms and technology stacks are not fully standardised. We therefore interviewed a number of companies about their approach to implementing smart contracts and distributed ledgers to tackle supply chain challenges. Our focus on the supply chain use case is not intended as an endorsement of the concept or of the companies interviewed; rather, we use the supply chain use case as a means of making the concepts discussed more concrete and tangible.

Use case: supply chains

Global supply chains delivering food, clothing, raw materials, consumer goods and many other things to businesses and consumers are at the heart of the

[10] For more information about Blokur, see: music:)ally (2017), 'Blokur talks blockchain music: The technology on its own is not the whole picture', http://musically.com/2017/08/03/blokur-blockchain-music-technology/

modern global economy. While every supply chain has its own unique realities and characteristics, they are broadly defined by their complexity and opacity. Both factors make it difficult for those involved - producers, suppliers, processors, distributors, retailers, consumers, etc - to trust the other parties in that chain.

Whether regional, intranational, international or global, supply chains often involve numerous actors with myriad different standards, quality assurance mechanisms and record-keeping processes. In addition, members of a supply chain often have to coordinate with various governments and third-party organisations to cross borders, pay customs or certify the authenticity of products.

To take a short example, even the production of a something as seemingly simply as a Christmas jumper involves numerous different parties and organisations. Before a shopper can purchase a new Rudolph jumper, the wool needs to sourced by a shepherd, shipped to a mill, processed into yarn, shipped to a knitting company or clothing manufacturer, woven into a jumper, shipped to a distributor, and finally shipped to a retailer. Every time the wool or jumper changes hands the interaction needs to be documented on a ledger with information such as time, weight, temperature, value, customs information or ownership rights.

If that jumper were to be shipped internationally or if it were to include other materials or design elements - eg an electric nose or plastic sequins - then the supply chain would grow even more complex and opaque. Given the varying processes and record-keeping standards of each link in that chain, the prospect of tracking the constituent parts of that jumper from source to finished product can be daunting. Complicating matters is the fact that often an individual member of a supply chain will have little to no contact with, or knowledge of, other members beyond those with whom they interact directly.

While many existing, centralised supply chains function adequately and effectively using traditional methods, there is growing concern about the lack of accountability and the lack of ability to trace the movement of goods through different supply chains. This is not only a concern for end consumers but, increasingly, for the businesses that rely on, and participate in, these supply chains. Issues such as unsustainability, lack of accountability, counterfeiting, poor working conditions or child labour are not only bad for society but can have a serious impact on the reputation of the businesses providing the end product - even when they are not aware or directly responsible.[11]

There are many who hope that creating a distributed, immutable, transparent ledger to record supply chain events, transactions and interactions might help tackle some of these issues. They argue that distributed ledgers may be able to create a trusted environment that can bring together the various parties involved in a supply chain and help to solve some of the challenges raised by the increasing complexity and opacity of global supply chains.

[11] For more information, see: Arthur Neslen (2017), 'Pepsico, Unilever and Nestlé accused of complicity in illegal rainforest destruction', https://www.theguardian.com/environment/2017/jul/21/pepsico-unilever-and-nestle-accused-of-complicity-in-illegal-rainforest-destruction

> *Supply chains, the way they're governed and how data is shared is so fundamentally broken that I think supply chains have been looking for a technology solution for a while. For a while everyone thought it was the 'Internet of Things' that was going to save the day, and now everyone has shifted into blockchain is going to save the day. It's a bit of both, but I think blockchains have shown people that work in supply chains this fantastic new frontier for collaborating and reinforcing rules.*
>
> ***– Jessi Baker, Provenance***

Supply chain case studies: Everledger and arc-net

Everledger[12]
Everledger tracks the provenance of high value assets on a distributed ledger. Using a combination of public and private blockchains, Everledger can give network participants the capability of viewing the ownership and transaction history of a high value asset, whilst still prohibiting access to sensitive or confidential information to those without permission.

Verifying the provenance of a high value asset is difficult, in part because a store-bought asset will have one document that proves authenticity and another that proves ownership. Keeping those documents associated with the asset is challenging, especially if one of them is lost or stolen. In Everledger's solution, metadata points are extracted from the asset to create a digital thumbprint - a record on the blockchain that is created using the asset's defining characteristics, such as its history, current ownership and dimensions. This thumbprint can then be used to verify authenticity.

In 2015, Everledger started working with manufacturers and retailers from across the diamond supply chain in order to create a historical ledger that documents the movement of diamonds. The ledger contains data about the origin and certification of the diamonds and processes involved in manufacturing them.

arc-net
By working with companies in the food and beverage industry to establish a chain of custody around their products, arc-net helps companies to improve their approach to product and supply chain authentication.

arc-net uses private blockchains within production environments that are designed to integrate with public chains if necessary. When products are first registered to the blockchain, they are assigned a Universal Unique Identifier (UUI) that references data about the product such as location of origin and DNA (if the product includes animal products). The product is mapped and tracked through processing, so when the final product is packaged it

[12] For more information about Everledger, see: https://www.everledger.io

receives a unique QR code that can be scanned using a smartphone to reveal the history and provenance of the product.

arc-net recently partnered with Adelphi's Ardnamurchan Distillery on a limited edition spirit production by marking and authenticating each bottle that they produced.[13] The bottles could be tracked from the distillery, through the network of suppliers and finally to the consumer. The consumer received confirmation and blockchain data for each unique bottle, thereby eliminating the possibility that the bottle was a counterfeit.

Because distributed ledgers produce a single, immutable, decentralised record of events and transactions, distributed ledger technologies have the potential to improve transparency and inform trust between members of industries beyond supply chains and the finance sector. As companies continue to work to implement these technologies over the coming years, it is likely that they will prove their value even more widely. Efforts at implementing distributed ledgers should focus on use cases or industries that could benefit from greater trust and transparency. In other words, businesses that are thinking about implementing distributed ledgers should understand what distributed ledger technologies can and cannot offer, and identify how distributed ledger technologies can be used to address their current real-world problems.

The remaining sections of Part 1 detail the emergence of smart contracts as an extension of distributed ledger technologies, and explore the role they might play in underpinning trust and creating efficiencies in interactions between people, businesses, governments and other organisations.

What are smart contracts?

The implementations of distributed ledgers discussed in the previous section have potential for underpinning trust in interactions between people, businesses and governments. However, perhaps some of the greatest promise of distributed ledgers lies in their ability to host smart contracts - sometimes referred to as the third generation of distributed ledger technologies.[14] There are many different definitions of smart contracts, but our preferred definition comes from Gideon Greenspan:[15]

> *A smart contract is a piece of code that is stored on a blockchain, triggered by blockchain transactions and which reads and writes data in that blockchain's database.*
>
> ***– Gideon Greenspan***

[13] For more information about arc-net and their distillery project, see: Charlie Taylor (2017), 'Arc-net and Scottish distillery in blockchain link-up', https://www.irishtimes.com/business/technology/arc-net-and-scottish-distillery-in-blockchain-link-up-1.3244056

[14] For further details of the 'generations' of blockchain technology, see our previous report: https://theodi.org/article/applying-blockchain-technology-in-global-data-infrastructure/

[15] Gideon Greenspan (2016), 'Why Many Smart Contract Use Cases Are Simply Impossible', https://www.coindesk.com/three-smart-contract-misconceptions/

This simple definition captures what most consider to be the essence of smart contracts; at their core, smart contracts are executable computer code stored on a blockchain or other distributed ledger. The aim of smart contracts is to enable certain actions to occur automatically within distributed ledger systems. Similar to most mature database systems, you can write out a set of rules within the database system that, when triggered by certain conditions, result in a change to the state of the database.

By providing a means to automate processes taking place within distributed ledger systems, smart contracts potentially offer a way of codifying more complex interactions than an immediate value exchange, all while retaining the benefit of informing trust inherent to distributed ledgers. For example, a business could programme a smart contract so that when a customer purchases a service, the smart contract is triggered and they are automatically sent a confirmation email and given access to the service.

Key features of smart contracts

We have identified seven key features that useful smart contracts should have. Specifically, we have focused on the features that need to be considered when trying to implement smart contracts to solve real-world problems.

- **Automated** – the execution of smart contract code requires no manual intervention
- **Deterministic** – given the same initial conditions, the executed code should always give the same result
- **Virtual** – both the conditions and the consequences of the contract must be represented within the ledger
- **Unalterable** – the conditions and consequences of the contract cannot be changed, except in ways that are originally anticipated within the contract itself
- **Irreversible** – transactions that occur according to the terms of the contract create permanent changes within the ledger
- **Available** – anyone (who has permission) is able to trigger the execution of the contract code at any time
- **Auditable** – the code, input and output of the smart contract is reviewable by any member of the network

The reason smart contracts might be used, as opposed to automation on a standard centralised system or database, is that they explicitly interact with a distributed ledger. This means that they are triggered by changes to the distributed ledger and make changes to the same ledger, and as a result can automate processes within distributed ledger systems.

They are also stored on the ledger, as opposed to being automated processes that just read and write to the ledger, meaning that the executable code itself has the same properties as other records on a distributed ledger. From these properties, especially immutability and transparency, smart contracts derive their potential to underpin and inform trust. Being stored on the ledger itself means that the person or organisation triggering a smart contract does not need to rely on the organisation or person who wrote the automated process to deliver the expected outcome, rather the system itself guarantees delivery, without intervention from any other party.

Smart contracts – misleading terminology?

The term ‘smart contracts’ is generally considered to be somewhat misleading. When asked to define a smart contract, many of our interviewees replied that most were neither ‘smart’, in the sense of exhibiting adaptive behaviour, nor ‘contracts’, in the sense of a legal agreement. While there have been a number of efforts to create legally enforceable smart contracts, use of this term is often held to be misleading given that most smart contracts are not intended to have a legal function.

“It has almost nothing to do with legal contracts at all – if the word contract was not used, no one would assume any connection to the legal context, so this is a source of gigantic confusion.”

– Interview participant[16]

There is clear potential for smart contracts to be used to automate some aspects of legal agreements, but usage of the term ‘smart contract’ should not be allowed to crowd out other potential implementations beyond this narrow definition.

Trusted interactions through smart contracts

The promise of smart contracts therefore lies in their potential to enable transparent, auditable and efficient interactions between people, businesses and governments. For a long time, automating parts of interactions has been seen as a source of efficiency and savings for businesses and other organisations. A rudimentary example might be an automatic monthly debit where a person or business agrees to let a service provider automatically debit payment from a given account.

In this example, this simple interaction is automated in such a way that makes it more efficient than it would be if handled manually. However, both parties must rely on a third party - the bank - to execute the arrangement, and trust it to do so in the way they both expect. While it is unlikely that a lack of trust would undermine such a simple arrangement, in more complex scenarios, relying on third parties can create risks, central points of failure and inefficiencies. These challenges arise not only from the complexity of the interaction - for example, payment for services based on a variety of criteria - but also the complexity of the networks involved, for example where a process relies on interactions between a wide range of actors who are not in direct contact.

It is in the more complex interactions and networks that characterise many real-world industries that many of our interviewees saw potential for smart contracts. Trust tends to be more dispersed in industries where networks of businesses and customers are dispersed by geography, role, size and other factors. In these industries, trusted third parties are less likely to exist and it can be difficult to cultivate trust across organisational and institutional boundaries.

Using smart contracts, businesses and other organisations involved in these complex networks are able to codify interactions directly into a shared, distributed, immutable database. By building on the properties of distributed ledgers, this enables automation of processes in a transparent manner, meaning that all actors in a given network witness the interaction. Not only can each member of the network review the behaviour codified into particular smart contracts, the interactions create a permanent, auditable record on the distributed ledger.

This transparency of interactions and the creation of an immutable record have the potential to remove the need for a trusted central authority or marketplace to approve, administer and record the interactions. Instead, these functions are carried

[16] This interviewee could not be reached to confirm attribution.

out by every member of the network in a distributed manner, engendering trust throughout the network.

Supply chain use case: smart contracts

Because supply chains are made up of a number of disparate companies with their own needs and aims, it is often difficult for the individual members of a supply chain to trust each of the other participants in that chain. Distributed ledgers and smart contracts offer the opportunity to increase transparency and engender trust between the various members of a supply chain and hold the potential to realise efficiencies in the interactions between those parties. It is because of this that supply chains are the most prominent use case for distributed ledgers and smart contracts.

Considering the important role supply chains play in society, any efficiencies and improvements that can be made have potentially huge benefits. Likewise, any errors and mistakes can cost businesses and consumers dearly. More than just offering a means to create transparency and auditability, many hope that smart contracts can help bring about efficiencies and confidence in the day-to-day management of supply chains.

> *All the kinds of problems that certification and audit aim to address could potentially be remedied or streamlined a huge amount by blockchains and smart contracts, and I think that's why everyone is getting very excited about it.*
>
> ***– Jessi Baker, Provenance***

Supply chain case studies: Provenance and Sweetbridge

Provenance

Provenance is a platform that enables greater transparency and traceability within supply chains. Provenance was one of the first companies in the world to use blockchain technology in supply chains. Their case study, 'From shore to plate: Tracking tuna on the blockchain',[17] demonstrates their approach to supply chain transparency and traceability from the origin of a product to the point of sale.

At the source of origin, Provenance worked with local fishermen to capture data about the collection of fish, using SMS messages to register each catch. This would issue a new asset ID on the blockchain for the product, containing data about its capture. The assets were then sold and transferred both physically and on the digital register to the supplier and could be identified using unique identifiers, such as QR codes. The history

[17] For more information, see: Provenance (2016), 'From shore to plate: Tracking tuna on the blockchain', https://www.provenance.org/tracking-tuna-on-the-blockchain

of the digital asset could then be verified by local NGOs using Provenance's blockchain explorer software. The NGOs would check the asset history against a recognised standard - eg if the fisherman is a member of the Pole and Line Foundation Association - to make sure that the fish have been sourced legally and ethically. From there, suppliers would register the fish, again on the blockchain. The digital asset would be verified by the trusted NGOs once more and this verification would execute a smart contract. The smart contract would signal that the fish was ready for processing and would allocate new digital asset IDs to containers, in which portions of the fish would be packaged, and link these new asset IDs to the original asset ID of the fish. Retailers who received the packaged products could identify each product using scannable shipping labels. The retailer would then attach an NFC smart sticker to each product, so that consumers could scan the sticker using Provenance's Item Tracking Interface and see the journey of the item.

The study aimed to provide proof of compliance to standards along the supply chain and determine whether distributed ledger technology could provide an open platform that would increase traceability and transparency of products in supply chains.

Sweetbridge

Sweetbridge are a blockchain-based protocol stack that works to enable supply chains to be more efficient without the involvement of third parties. They aim to enable companies who exist in a shared ecosystem to communicate with each other and share assets, and to make it easier for assets to be accurately valued by members of the ecosystem. Having these communication links and being able to accurately put a value to assets is particularly important for smaller companies, who may want to have access to several public ledgers in order to work with multiple larger organisations, but also need assurances that assets retain their specific value as they make their way from the start of the supply chain to the end.

Sweetbridge place an emphasis on the importance of building applications that increase the liquidity of financing and payments while enabling the efficient flow of product and state information on top of distributed ledger technologies. Smart contracts – which Sweetbridge refer to as 'programmable work flows' – can be one of these applications, and are usually introduced in order to implement logical processes, such as automatically allocating tasks to different companies in the supply chain in response to the record of a new transaction.

The companies that Sweetbridge work with often want tailored, permissioned applications on a private blockchain that help them maintain their competitive edge, but also enable them to foster their own ecosystem within their specific supply chain by giving permissioned access to close business partners, such as suppliers and retailers. This type of application can be beneficial to most of the actors within the ecosystem, although much of the control remains in the hands of the company for whom the application has been built.

Are smart contracts and distributed ledgers right for my business?

Whether or not businesses choose to implement smart contracts depends considerably on the specifics of the potential use case, the industry context and the needs of those involved. Because of this, and the fact that the technology is still emerging, there are few established rules or parameters that can be universally applied when deciding whether to pursue smart contracts.

One widely offered piece of advice expressed by our interviewees was that businesses should make sure that they identify a problem first, before considering the application of these technologies. The problem businesses are trying to solve should be well-defined and characterised by a lack of trust between the actors involved. In addition, it should exist in an industry or context where there is no central authority that would be trusted by those involved to administer a centralised solution.[18]

> *For people who are 'outside of Blockchain', who are coming from the corporate world, who are doing various other kinds of things - they are approaching it from first principles and they're approaching it from what are their business problems. What they're seeing is there's a technology that enables them to attack old problems in new ways.*
>
> ***– Interview participant[19]***

While the vision and promise seems clear in use cases that fit these general criteria, an additional question businesses should ask is whether smart contracts and distributed ledgers can deliver on their promise today, in tackling real-world problems, without requiring an unprecedented industry shift in business practice.

> *We have some partners and clients who are really bought into the fully trustless method and believe that might be the only way in order to actually create transparency and traceability, and I agree that the only way we'll ever get full traceability and transparency in the world's supply chains is through something that is fully decentralised and trustless. But I think we've got a long way to go.*
>
> ***– Jessi Baker, Provenance***

In many cases, it will be the needs and desires of a company's clients and the industries in which that company works that will define whether or not smart contracts will prove useful. By understanding that industry context and the general criteria around lack of trust, businesses should be able to make reasonable judgements on the applicability and potential usefulness of these technologies.

[18] For more information, see: Gideon Greenspan (2015), 'Avoiding the pointless blockchain project', https://www.multichain.com/blog/2015/11/avoiding-pointless-blockchain-project/

[19] This interviewee could not be reached to confirm attribution.

> *It's important for us to make sure that we solve a problem that our clients and industry have so we need to find the best way to solve that problem. In order to do that you need to have a pragmatic approach to say the problem is happening today in today's world, so you need to understand what the best option is today.*
>
> ***– Calogero Scibetta, Everledger***

Part 2: Informing trust through traditional v high-tech mechanisms

Businesses that identify potential uses for distributed ledgers and smart contracts in enabling trust will need to decide how to approach three key challenges when implementing new systems.

Through our research, we have identified three key challenges that businesses implementing systems based on smart contracts should consider. Each of the three sections in Part 2 presents a challenge or set of challenges that might affect the design and performance of a system using smart contracts and distributed ledgers. We work through each of the challenges in the context of the supply chain use case and examine potential approaches to tackling these challenges.

Within each of these approaches there are choices that have to be made - choices that revolve primarily around the need to choose between using existing trust mechanisms to guarantee parts of the system or designing an approach that draws upon the technical capabilities of the system. For each, the choice is unlikely to be binary and will often depend on the context of the particular use case, meaning that businesses will have to navigate their way between the ideal and the pragmatic - between what might be possible given an ideal setting, and what is practical and achievable given a particular industrial, regulatory, or social context.

Challenge 1: Representing the real world

> *It's important to make a distinction between the perfect world - where everything exists in a blockchain and everything can be resolved through consensus - and the real world.*
>
> ***– Jess Houlgrave, Codex Protocol***

Every application of distributed ledgers requires some form of data to be recorded into the database. In almost all cases beyond cryptocurrencies, the reason for recording this data into the ledger is to create an immutable record of some real-world event or interaction. This raises a fundamental question of how within a distributed ledger one can represent data about a thing or event in the real world in a way that guarantees the data being stored in that digital ledger is in fact an accurate representation of that real-world thing or event. Put another way, to what extent can distributed ledgers, as a digital technology, guarantee the truth of the physical world?

One person we interviewed argued that because of the origins of blockchain, many people remain confused about what the technology is capable of doing in terms of guaranteeing data integrity.

> *Bitcoin causes people to get very confused about this kind of thing because Bitcoin has this miraculous property of basically existing entirely inside the computer network.*
>
> ***– Interview participant[20]***

Thibaut Schaeffer from Provenance perfectly encapsulated this misconception and the problems that stem from it: "Obviously blockchain gives you really strong guarantees on what happens to data once it's in the system but it doesn't give you any guarantee about what the data is when it enters the system." In short, if you're going to trust the data stored in a distributed ledger, you need to work out mechanisms for injecting data into the ledger in a trusted manner.

In addition to this, smart contracts themselves cannot access off-chain data sources directly. This is due to the distributed nature of blockchains: if a decentralised smart contract were to call a web-based API when the conditions within the ledger were met for nodes to execute the code, it would trigger this API call from every computer or node on the network. This would likely have a huge impact on the infrastructure of the data source (there were over 17,000 nodes on the public Ethereum network at the time of writing.[21])

Further complicating matters is the fact that if the data being accessed by a smart contract is prone to fluctuations – for instance, real-time market prices or temperature readings – different nodes of a system might receive different data in response to the same API call, even if the calls are only separated by a matter of seconds. This would cause massive issues within the system, since smart contracts run by different nodes could come to different conclusions and therefore have conflicting copies of the database. Were this to happen, it would be more difficult for the network to reach consensus on which contracts to execute or what the next version of the database should be. Ultimately, it would hinder the system's ability to guarantee a single version of truth or enable trust between different actors in the network.

Supply chain use case: impact of weather on logistics

By deploying distributed ledgers and smart contracts within a supply chain it would be possible to create a system wherein a supplier could trigger a smart contract designed to record the impact of poor weather conditions on the transport of goods, thereby creating a verifiable record that a certain supplier was not responsible for the late arrival of a shipment of goods. That smart contract might be designed to verify the data input by the supplier by comparing it to weather readings taken by a national weather service such as the UK Met Office. The contract would initiate an automated call to the Met Office API and, depending on whether the response aligned with the data input by the supplier, would then create a record validating the supplier's claim.

The challenge in implementing such a system is that since smart contracts need to be deterministic, each node that executes the contract code must come to the same conclusion. If every node has to call the Met Office API, it

[20] This interviewee could not be reached to confirm attribution.
[21] For current network size, see: https://www.ethernodes.org/network/1

might overload the Met Office system, meaning some nodes do not get any data and therefore can't execute the contract correctly. Furthermore, if a node that received an initial null response attempts to execute the contract again, it may not get the same weather reading from the API, leading to the creation of conflicting records within the system. For example, a contract could be written that would execute if the temperature in a certain location dropped below 0 degrees Celsius. If, after half the nodes called the API, the temperature reading was updated to above 0 degrees, the nodes would come to different conclusions when running the contract.

As a result of these complications, data about real-world events must be injected into the ledger *before* any smart contracts attempt to execute their terms. Only after data has been recorded into the ledger can smart contracts be allowed to use that data to evaluate and trigger the contract. In the supply chain example above, this would involve injecting temperature readings from the UK Met Office API into the ledger, then allowing the supply chain smart contracts to consult the stored temperature value. This allows smart contracts to perform their operations entirely within the copy of the database held by each node, such that each node will receive exactly the same temperature reading and result in exactly the same decision about whether or not to execute the terms of the contract. This has the added benefit of creating an auditable record of the data being used - a record that could be used to understand the behaviour of a smart contract after the fact.

The challenge that this presents for businesses that have decided to pursue smart contract implementations is how to record data in the ledger in a way that all members of the network can trust the accuracy of that data. This challenge exists even where the writing or reading of data from the ledger is handled automatically off the ledger, for example where an 'Internet of Things' sensor records objective measures, such as temperature or location,[22] or where a smart contract unlocks a 'smart lock'.[23] This is because the distributed ledger cannot guarantee that the automatic, off-chain sensor will always behave in the expected manner. Even in fully automated systems, power cuts, faults, connectivity outages or vandalism can break the connection between the physical and digital world.

Recording trusted data into the ledger using oracles

The most common way of recording data into a distributed ledger in a trusted manner is through 'oracles'. Oracles are people or businesses that are tasked with recording specific real-world data into the distributed ledger. In some cases oracles can even be tasked with reading the ledger and putting something into action in the physical world.[24] The presence of an oracle within a system to some extent reintroduces the need to place trust in third parties: the oracles are *trusted* to provide accurate data. However, oracles do not necessarily reintroduce the need to trust a single third party, since it is possible for multiple oracles, performing multiple functions, to exist within a single system. For example, the existence of multiple

[22] For more information, see: Anna Hensel (2018), 'Nevada's Filament unveils chip that lets industrial IoT devices communicate with blockchains', https://venturebeat.com/2018/01/16/nevadas-filament-unveils-chip-that-lets-industrial-iot-devices-to-communicate-with-blockchains/

[23] For more information, see: Giulio Prisco (2015), 'Slock.it to Introduce Smart Locks Linked to Smart Ethereum Contracts, Decentralize the Sharing Economy', https://bitcoinmagazine.com/articles/slock-it-to-introduce-smart-locks-linked-to-smart-ethereum-contracts-decentralize-the-sharing-economy-1446746719/

[24] For more information, see: Stefan Thomas, Evan Schwartz (2014), 'Smart Oracles: A Simple, Powerful Approach to Smart Contracts', https://github.com/codius/codius/wiki/Smart-Oracles:-A-Simple,-Powerful-Approach-to-Smart-Contracts

oracles providing similar types of data allows businesses to choose between trusted authorities, whereas in other cases trusted oracles may provide authoritative sources of different types of cross-referenceable data. Either of these ultimately results in a wider distribution of trust within the system when compared to a system where a single authority provides or collates all different types of data.

If oracles are required for a particular use case, businesses will need to inform and incentivise trust in those oracles to provide accurate data against which smart contracts can be run. From our interviews, there were broadly two approaches taken by companies building systems based on distributed ledgers.

i) Informing trust through existing, traditional trust mechanisms

Several of the companies we spoke to were designing systems where data would be injected using existing trust mechanisms in some way. In particular, there were several cases where our interviewees did not see an issue with using an official organisation or authoritative body to act as a trusted third-party oracle. In certain cases, it might be possible to use trusted sources of data - or versions of official truth - that are publically available, potentially as open data. However, in other cases, the members of a network may wish to directly involve an authoritative body in inserting the required data into the ledger - though this would likely require some incentivisation of the authoritative body. Incentivising participation might be useful in guaranteeing that the trusted organisation continues to provide the required data over the lifetime of a smart contract.

Another case where traditional trust mechanisms may have a role to play in inserting data into a distributed ledger would be when an organisation is responsible for generating data based on their reputation and expertise - for example, organisations tasked with certification.

In cases where trusted third parties are required, there are often existing non-technology-based incentives and mechanisms to enforce their behaviour, such as the desire of an authoritative body or trusted organisation to maintain their reputation or standing. While these types of systems place trust in third parties, many of the companies developing such systems believe they still improve efficiency and are preferable to systems that do not utilise distributed ledgers.

Supply chain use case: certification

One of the main use cases for trusted third-party oracles involves certification – where an organisation or initiative creates a set of standards for supply chain behaviour, audits supply chain businesses against these standards and awards businesses with a ‘certificate’ if they comply with these standards. These authoritative, neutral certification bodies often focus on particular supply chain issues in certain industries, such as the Marine Stewardship Council (MSC), which focuses on sustainability in the fishing industry.[25]

These types of organisations could potentially act as trusted oracles, using distributed ledgers to issue certificates for goods that meet predefined standards. These could help businesses to provide a transparent and immutable record of certification to consumers and other businesses. As a result, every member of the network would be able to verify that the certification came from a reputable certification body, thereby circumventing the need to trust the word of the supplier.

Furthermore, certification could be carried out through smart contracts, where suppliers could submit evidence of compliance through the system

[25] For more information about the MSC, see: https://www.msc.org

in a transparent manner. Similarly, supply chain smart contracts could be triggered by the process of certification. For example, a purchaser's smart contract could be set up to only execute if the goods in question have been certified by one of these certification bodies. These applications of smart contracts could improve the efficiency with which certification or proof of certification takes place.

> *"Having a centralised oracle is definitely one way and I think it works for many, many cases – especially for certification – where you want to know that someone is certified by a particular entity."*
>
> **– Thibaut Schaeffer, Provenance**

ii) Informing trust by drawing on the technical capabilities of the system

While in some cases businesses implementing smart contracts are relying on existing trust mechanisms to engender trust in the data inserted into the ledger, other businesses are pursuing ways of using economic incentives and the technical capabilities of distributed ledgers to engender trust in that data. In moving away from a system that relies on trusted organisations to inject data, businesses are attempting to get closer to a system that internally rewards the provision of accurate data. As Thibaut from Provenance explained:

> *For other things, you have to move away from the centralised expert to some extent, and blockchains provide ways to engineer governance directly using the currency mechanisms that you have built into blockchain. You can incentivise data submission from the crowd and there's not this problem that you have to trust a third party.*

The idea of incentivising data submission using economic incentives built into a distributed ledger system is at the heart of the system being developed by Oraclize.

> *You don't want to just trust our third party. The third party should not be trusted. What we do is we prove through some cryptographic techniques that the data we fetch was not tampered with or altered by us. We provide a safe layer, which can prove to you that we didn't touch the data, and this safe layer comes with the trust of a much bigger party. Perhaps it could be Amazon or Google or Intel or Microsoft. It's very difficult not to have the initial trust party but the idea is that if you can have many large and competing third parties then you have much higher*

guarantees of security because to tamper with the data they will have to all cooperate, and even then they could be found out.

– Marco Giglio, Oraclize

Deciding how to guarantee the integrity of data

All companies looking to implement smart contracts on distributed ledger systems to solve real-world problems are likely to have to access data from beyond the system itself. Businesses will therefore need to decide how they will go about getting real-world data into the ledger in a trusted way and, importantly, how they will guarantee the continued functioning of any system they develop. The method businesses choose will depend on the type of data involved, the existing context of the industry, and what exactly they require from the system being built. Indeed, there are likely to be lots of use cases where both approaches might be required for different parts of the system.

Supply chain use case: example decisions

- If a smart contract requires weather data, it may be enough for all actors to trust an official weather service.
- If a company needs to know whether a contractor is certified, then it may be enough to trust data recorded by a certification authority.
- If a smart contract requires data about the market price of a product, say a Rudolph jumper, then it may be better to rely on incentive mechanisms and the technical features of distributed ledgers.
- A single system might need multiple approaches, for example requiring certification by central authorities but a marketplace approach for contextual data.

Relying on a trusted third party has the advantage of being simpler to implement since it builds upon existing trust mechanisms. However, the drawbacks are that it may to some extent recentralise trust within the system, and in some cases it might not be feasible to find sufficiently trustworthy third parties within a given system.

Relying on cryptographic proofs and economic incentives has the advantage of maintaining the distribution of trust in a way that should not be easy to manipulate or control. However, these approaches are currently unproven and might be difficult to design. Incentivising accurate data provision may require additional funding within the system and resolving multiple data sources might require more processing, making these mechanisms more expensive to run.

Tackling the challenge of representing the real world

Businesses looking to use distributed ledgers and smart contracts to solve real-world problems will need to:

- assess the existing levels of trust within their particular industry or sector, and the requirements of the network

- identify the most appropriate injection mechanism for each type of data that needs to be recorded in the ledger. Note that for different types of data, the appropriate mechanism may be different or may involve a mixture of various approaches. Businesses could:
    - consider using traditional trust mechanisms in cases where trusted third parties already exist
    - explore building incentive mechanisms to validate and record data in cases where there are no authoritative sources

Challenge 2: Edge cases, bugs and arbitration

> *Automation of processes where everything just runs like clockwork will be the case for 99 per cent of things. The areas where you would need human intervention are the edge cases, and the unwillingness to accept that there will be edge cases is potentially fatal.*
>
> ***– Interview participant[26]***

Ensuring that there are mechanisms to record accurate data into a distributed ledger in a trustworthy manner is a key part of enabling smart contracts to be used to automate a variety of processes. However, even where smart contracts execute based on accurate data, there will almost certainly be times where they do not function as expected. Despite the high-profile cases of malicious attacks like the failure of the DAO[27] and the various vulnerabilities identified in smart contracts,[28] most of the people we interviewed for this report believed these failures represent edge cases rather than inherent flaws in smart contracts as a whole. These edge cases, however, present a challenge for any business looking to use smart contracts to tackle real-world issues.

The challenge of developing smart contracts

Such edge cases may arise from foreseeable and unforeseeable changing real-world conditions that invalidate their function, such as a supplier in the supply chain going out of business. However, the majority of reported cases of smart contracts not functioning as intended have arisen from malicious or unintended exploitation of flaws in smart contract code.[29]

One often cited reason for these vulnerabilities is the fact that the languages and environments that smart contracts are written and operate in are relatively new, and are in many cases still under development. As with any emerging software environment, this is bound to lead to edge cases where systems behave in unexpected ways since the systems themselves are adapted through use and the best way of implementing a given task is often defined ad hoc by individual programmers. In line with expectations from software development, these issues should decrease over time. However, the speed at which smart contract-based applications are being pushed live, and the immutability of smart contracts once they are written into a distributed ledger, means any potential vulnerabilities or flaws may have an impact on real-world implementations for longer than usual.

It can even be argued that the very same properties and features that make smart contracts potentially useful for automating business processes in complex industrial environments ultimately make them more vulnerable to exploitation than existing methods of automation. For example, because smart contract code stored on a

---

[26] This interviewee could not be reached to confirm attribution.
[27] David Siegel (2016), 'Understanding The DAO Attack', https://www.coindesk.com/understanding-dao-hack-journalists/
[28] Ivica Nikolic, Aashish Kolluri, Ilya Sergey, Prateek Saxena, Aquinas Hobor (2018), 'Finding The Greedy, Prodigal, and Suicidal Contracts at Scale', https://arxiv.org/abs/1802.06038
[29] Ivica Nikolic, Aashish Kolluri, Ilya Sergey, Prateek Saxena, Aquinas Hobor (2018), https://arxiv.org/abs/1802.06038

distributed ledger is immutable, it cannot be directly altered, updated or patched in response to issues in the same way that almost all other software development can be. To deal with this, most smart contracts are created with a kill switch – a way of stopping them from carrying out their terms – which is triggered when a new version goes live. However, these switches have also been exploited to carry out 'unauthorised' shutdowns of functionality.[30]

Another factor that contributes to the vulnerability of smart contracts is that they are difficult to test before deployment, especially where they interact with other contracts or real-world processes. Once deployed, it may also be harder to catch potential bugs before they are exploited since network members are not only able to trigger a smart contract immediately but concurrently in large numbers. Finally, all code is stored on the ledger, which every member of the network has a copy of; this might make the contracts easier to exploit, though lessons from open source software development indicate that having many eyes on code can equally make it easier to identify and fix bugs.

Challenges in the technical development of smart contracts

- Emerging languages and environments are naturally difficult to write code in
- Immutability means flawed code cannot be directly altered, updated or patched, meaning a new contract must be launched and the old one killed
- Kill switches designed to stop smart contracts from functioning can be maliciously or unintentionally triggered
- Complex behaviour and dependencies can be difficult to test before deployment
- Bugs can often be exploited at any time, by any member of the network with permission, until a contract is killed

While the problems caused by these potential vulnerabilities may seem insurmountable, many depend on context. Most of the recent exploitations of smart contracts have involved non-permissioned contracts on public ledgers. The types of malicious behaviour likely to exploit smart contracts may in fact be limited to these cases, especially when there are large financial incentives, in the form of cryptocurrencies, for carrying out exploitations.

In many potential applications of smart contracts, malicious exploitation is unlikely to be as highly incentivised for the participants – especially in cases where the network is primarily made up of businesses with reputations and relationships to maintain. While malicious exploitation might be less likely to occur in these cases, it is still more than likely that a not insignificant number of smart contracts will naturally contain bugs which cause them to behave in unpredictable or undesirable ways. In these cases, if a contract executes counter to the way it was intended to function, this discrepancy would need to be amended.

In cases where smart contracts execute in an unexpected manner, a mechanism for appeal should be in place if the system is to be trusted by its users. In many use cases, it might be possible for this to be resolved through correction by the entity responsible for creating that smart contract. In other cases, where the contract behaves in an unexpected manner, but the party who created it is unable or unwilling to correct this behaviour, some sort of arbitration mechanism will be required to settle the dispute. Finally, related to the first challenge of data integrity, smart

[30] For more information, see: https://github.com/paritytech/parity/issues/6995

contracts may trigger or respond to inaccurate data injected by an oracle – either deliberately or by accident. This would also require some form of arbitration.

Supply chain use case: dispute requirement examples

During the lifetime of a supply chain smart contract, such as the earlier weather example, the contract might at some point, for whatever reason, receive an extreme temperature or other input value. The result could be that the smart contract does not record or verify the suppliers' claim. In this situation, the supplier might want to dispute the outcome of the smart contract. It is possible that the party who created the smart contract would be able to issue a correction and rewrite and deploy the contract to avoid further bugs. However, they may be unwilling to do this, arguing that the oracle (the Met Office in this example) is responsible. In this case, it is likely that if the Met Office, as an oracle, is not willing to take responsibility, some form of dispute resolution will need to occur.

Resolving disputes

As with the first challenge of injecting data about the physical world into a digital ledger, there are broadly two different ways to approach this complication when it comes to smart contracts – either placing trust in traditional existing mechanisms to arbitrate disputes, or trying to create resolution mechanisms within the system itself using the properties of distributed ledgers and economic incentives.

As Thibaut from Provenance put it, within these emerging systems you "start seeing arbitration as a service", so "if there's a dispute, you can delegate the resolution of that dispute to a kind of 'arbitration court'." He went on to explain that you can run this court in different ways:

> *You could either run it through a trusted body so that would be like an actual court, or you could run it with economic incentives, for example by choosing people at random and rewarding them if they are honest and do their duty, make the right choice.*
>
> ***– Thibaut Schaeffer, Provenance***

Whether a company chooses the former or latter option will again depend on the use case, the existing conditions within the industry in question and what the company is designing the system to do. Since smart contracts will likely play a role in already-established systems, many of those systems will already have mechanisms for dispute resolution that could be used.

i) Existing trust mechanisms

Systems that rely on existing trusted third parties have the advantage of minimising the requirement for industries to adopt new governance regimes. In many use cases, existing legal mechanisms for dispute resolution already exist that would likely be able to fulfil this role. However, this would require legal professionals to understand the functioning of the smart contracts in place and for both parties to place trust in these legal systems.

Placing trust in these legal systems might be particularly challenging when the use case stretches across legal jurisdictions (as with global supply chains) since different members of the network might have different expectations regarding legal process, and different jurisdictions might have different interpretations of agreements. Perhaps most importantly, traditional dispute resolution mechanisms are precisely the types of systems that are currently deemed as inefficient and expensive, and therefore often what the adoption of smart contracts is intended to minimise.

ii) Distributed dispute resolution

For this reason, some businesses are exploring new ways of resolving disputes using the distributed nature of the systems they are building. Systems that rely on members of the network to act as arbiters have the advantage bypassing the need to trust a central administrator to define what counts as truth within the ledger, thereby avoiding the costs associated with traditional mechanisms. One approach currently being explored by Ocean Protocol, a company building a decentralised data exchange based on blockchain, would be to establish a formal reputation system within their platform to incentivise trustworthy behaviour.

Others are looking at how consensus governance mechanisms and distributed pools of arbiters could be used to resolve disputes and how economic incentives might be set up to reward just arbitration.[31] The challenge with such systems is that they may be difficult to establish, potentially add greater cost to the system and could in some cases mirror the inefficiencies of existing mechanisms.

Going off-chain?

In response to the challenges of edge cases, bugs and the need for arbitration, some have proposed that the solution is to store and execute smart contract code off the ledger itself. This would result in automation of processes that could read and write into the ledger but the execution of these processes would not be carried out in a distributed manner. While there are clear benefits to avoiding edge cases, this approach sacrifices some of the central proposed benefits of smart contracts – in particular, such a system would not be able to independently execute the terms of a contract. For this reason, there are those who argue this type of automation cannot be considered equivalent to the type of automation achieved through smart contracts.

> *"People that are doing a lot of the off-chain stuff, it's kind of square peg in a round hole type, if they're trying to call it blockchain when really, it's not. It's off the chain. It's not smart contract workflow driven, it's not being recorded to a large extent on the distributed record."*
>
> **– Todd Taylor, Supply Chain Advisor at Sweetbridge**

All companies looking to use smart contracts to tackle real-world problems will need to plan for edge cases, decide how they want to deal with them and determine how they will enable users to resolve disputes and issue corrections. The choice of resolution mechanism will be driven by the particular use case, the existing systems in place and the specific function the smart contracts are being asked to perform. In many cases, there might be the opportunity to use new technological approaches to resolve some forms of dispute while relying on existing, traditional mechanisms, such as the legal system, for complicated or especially contentious cases.

[31] For more information, see: https://confideal.io/

Tackling edge cases, bugs and arbitration

Companies working to implement distributed ledgers and smart contracts will need to:

- understand that while smart contracts can automate a number of different business operations, there will still be cases where things go wrong
- plan for edge cases and choose the most appropriate mechanisms for resolving disputes and issuing corrections when they arise. They could:
    - consider using traditional mechanisms in cases where existing arbitration would be sufficient
    - explore using the features of distributed ledgers to provide technological mechanisms for dispute resolution in cases where existing mechanisms may be cumbersome or not fit for purpose

Challenge 3: Cryptocurrencies and financial incentives

Another key aspect of distributed ledger systems that companies exploring how smart contracts might be used to tackle real-world challenges need to consider is how to pay for the system – both its implementation and operation. In a related way, one of the main challenges facing businesses will be how to incentivise other parties – other businesses, certification bodies, NGOs etc. – to use the system. Since the main potential benefits of distributed ledgers and smart contracts for underpinning trust come from the network, businesses must ensure these systems have the participation and trust of all the parties involved in order to create value.

Paying for the operation of a distributed ledger

Like all databases, distributed ledgers require digital storage and processing power to carry out operations within the database. As explained earlier in this report, public distributed ledgers use a variety of consensus mechanisms as part of their design – relying on financial incentives, in the form of cryptocurrency, to incentivise and reward participation in the process of creating new versions of the distributed ledger. While some systems use a communal pot to pay for some of this reward, most also use transaction fees from users to pay for these incentives – meaning anyone who wants to commit data to the database must pay to do so.

This requirement to provide financial incentives is increased in the case of smart contracts. As discussed earlier, when the required conditions are met, every computer executes the smart contract code independently to understand and verify the effect this operation has on the ledger. Every smart contract that is executed takes computing power to resolve, with more complicated contracts taking longer and using more resources to process. Because of this, a number of public systems charge a usage fee,[32] based on the complexity of the contract. This fee is payable by the person or company triggering the contract – and is required every time they trigger it.

Developing applications on public infrastructure has many benefits related to stability, integrity, transparency and independent verification. However, with more complex use cases, and volatile cryptocurrency prices, the cost of administering complex smart contracts may becoming increasingly prohibitive for businesses. In response, some companies developing smart contract applications have focused on permissioned and private ledgers – the main benefit being that there is no requirement to create and distribute a cryptocurrency with economic value, though these can be used if so desired. Because only a certain set of predefined organisations can run the nodes of a permissioned or private network, there is no need to incentivise unknown parties to host and administer the database.

This can have the impact of reducing the overall cost of operating the system, even when operating many complex smart contracts simultaneously. However, because private and permissioned ledgers are generally smaller and less transparent, they arguably sacrifice some of the guarantees of public ledgers related to integrity and stability. In addition, by controlling who has access to the system, the governance of the system and its infrastructure becomes, in a sense, centralised. This may limit the amount of trust that those outside the system are willing to place in the network itself.

Supply chain example: centralised setup and governance

Another challenge that arises when operating smart contract systems comes from who is paying to establish the system. Most of the companies creating supply chain applications that we talked to were working with major

[32] known as ‘gas’ in the case of Ethereum.

companies who are paying to establish the infrastructure to manage the supply chains they are involved in. Where systems are being developed and paid for by the major market player in a supply chain or ecosystem, it raises questions as to whether such a system will be able to realise the benefit of distributing trust since governance of that system will be centralised.

> *"In these private blockchains there's a primary entity and usually it's the big brand that's responsible for pulling everybody together and distributing product; that's taking orders primarily and then trickling those back out through the supply chain. They're the ones that are selecting and inviting, registering and permissioning the different participants and nodes on the network and then when things go awry, they're the ones that probably have the dashboard and the analytics in place that allow them to go back and query and find out what happened and see if they need to do some things to remedy an inappropriate action."*
>
> **– Todd Taylor, Sweetbridge**

The key question for businesses looking to implement smart contract systems is how to pay for the operation of the system in a way that retains trust of all the members of that system. This will likely depend on the purpose of the system, the level of trust between members of the network, how comfortable those members are with having their transaction records stored publically, and the willingness within the industry to contribute to such a system.

> *There are multiple commercial models. We define them based on the sector in which we operate. For example, in the export market it's actually the processor that pays for the service because the processor owns the supply chain, the farm and contract, the produce out. In terms of a retailer, in that case if it's specific to the retailer then the retailer pays.*
>
> ***– Kieran Kelly, arc-net***

Supply chain use case: operational choices examples

In a supply chain, a number of different approaches might be taken to paying for the setup and operation of the system:

- In a supply chain made up of similar-sized organisations who agree on an approach, the various members might choose to distribute the cost of setting up a private or permissioned ledger
- In a supply chain with only a few equally-sized organisations interested in a solution, each might chose to run a node and fund the cost of their own transactions in a public ledger
- In a supply chain with one large company, the central player might pay for someone to design and setup a private or permissioned system, which others can join at relatively low cost

Transacting value and incentivising participation

In addition to paying for the setup and operation of the platform, any company looking to use distributed ledgers and smart contracts to tackle real-world business challenges will need to decide whether the system will be used to process financial, monetary or value transactions. Since distributed ledgers were originally designed to facilitate cryptocurrency transactions, they can be used to process financial exchange within the system. Smart contracts provide a means to automate these exchanges of value, and this is one of the primary use cases within the finance industry – for example, Clearmatics are aiming to automate financial transactions between businesses.[33]

For some people, systems underpinned by the exchange of financial value through the system itself (using cryptocurrency or cryptocurrency-like systems) are the only systems that deliver on the promise of smart contracts – in particular in reducing friction in business processes.

> *The digital currency gives them an opportunity to settle much more quickly, to recognise payment much more quickly, to distribute those payments much more quickly across their value chain and to provide financing to some of the strategic partners within their value chain that might be paying obscene amounts of interest today that are hurting the value chain overall.*
>
> ***– Todd Taylor, Sweetbridge***

By handling payment within the system, the blockchain can guarantee the integrity of payment and, using the features inherent to distributed ledgers, can encourage good behaviour through financial incentives.

> *One way we can use incentives is to ensure that people are doing the right thing when they are using the system. In the sense that if things are free, people tend to abuse them. Very small transactions fees can often help to create the right incentives.*
>
> ***– Jess Houlgrave, Codex Protocol***

[33] For more information about Clearmatics, see: https://www.clearmatics.com/

Supply chain use case: exchanging value through a distributed ledger

In a supply chain, transactions like paying for a shipment of goods could be administered within the system using a cryptocurrency or cryptocurrency-like system. Handling the transactions in this way within the supply chain would provide incentives to participate, potentially enabling faster transactions and payment clearing. In addition, it might also enable businesses within the supply chain to automate some of these interactions in ways that would not be feasible in centralised systems.

For example, a smart contract could allow the automatic release of payment in response to the injection of data from a trusted third party confirming the delivery of a specified product. Both the buyer and seller in this example could trust the system to carry out the terms of the contract – even if neither completely trusted the other party. Finally, the inclusion of transaction costs might encourage members of the network to ensure they are using the system effectively and efficiently.

However, many businesses trying to build applications on top of distributed ledgers have concerns about building cryptographic asset exchange mechanisms into their products and services. Some of the concerns raised during our interviews stem from the needs of users of the system, both in terms of usability and comfort. Many of our interviewees questioned whether users would be willing to learn the procedures required to actually use cryptocurrencies and questioned whether they would be prepared to rely on existing volatile cryptocurrencies as a mechanism for exchange. In response to this, some businesses have taken the approach of using the distributed ledger to keep a record of payments, but actually process those payments off the chain using traditional methods, such as bank transfers.

> *We have never been involved in cryptocurrency at all. We don't see a cryptocurrency element as a must have because you will have specific transactions and needs and these will always occur in fiat currencies but we don't necessarily need crypto for that. For me it's not a must - you can if you want, but…*
>
> ***– Calogero Scibetta, Everledger***

Although these types of approaches might be accused of sacrificing some of the advantages of smart contracts where payment is handled entirely on a distributed ledger, some argue that the systems retain advantages over manual, centralised approaches.

> *It's certainly a misconception of lots of people that making payments happen faster is exclusively to do with the payment system and that's not really true. We can send money quite quickly to lots of different places; the reason why people get paid late is nothing to do with whether you're using Bitcoin or Ether or PayPal or a bank or whatever, it's to do with working out who is the right person to pay and all of the challenges that go into all of that. You can separate out the two things and so I concluded some time ago that getting people paid faster and more reliably did not require using cryptocurrency.*
>
> ***– Phil Barry, Blokur***

Whether or not systems require or enable users to exchange value through the system does not necessarily have an impact on how the system is designed to run. For example, systems that do not enable value exchange can be built and run on existing public ledgers, and only the operation of the system needs to be paid for in cryptocurrency. However, this is not exposed to the users of the system.

As with both data recording and edge cases (the first and second challenges outlined above), it is possible to approach these payments in different ways – by relying on distributed ledger-based cryptocurrencies to process transactions, relying on traditional payment mechanisms, or not recording or transacting value at all. Again, these approaches are not necessarily exclusive; within a single platform designed for a particular use case, different parts of the system could adopt different approaches.

> *As far as cryptocurrency, that is just one aspect of the technology. Codex is an infrastructure layer and not a direct application layer, so applications are built on top of that infrastructure - some of which may involve individuals interacting with cryptocurrency and some of which won't.*
>
> ***– Jess Houlgrave, Codex Protocol***

Whether or not businesses decide to exchange value through the distributed ledger will depend on the industry context and the willingness of potential users to participate. Our interviewees broadly agreed that exchanging value through the system itself might be useful, but that it was probably out of reach in most current use cases.

> *It's going to take a stable token, a stable coin and that's not easy to create. Being a central bank is not an easy job, but settlement can be done, and trade network ledger can be done and it's something that we're working on.*
>
> ***– Jason English, VP Protocol Marketing at Sweetbridge***

Tackling the challenge of cryptocurrencies and financial incentives

For those companies that choose to implement distributed ledgers and smart contracts to tackle real-world problems, it will be important to:

- decide how the system will be governed and who will pay for its operation. The answer to these questions will depend on the use case, the existing systems in place within the industry in question, and the needs and desires of the parties involved:
    - In certain cases, businesses may prefer to govern the system cooperatively and pool their resources to pay for the operation of the system
    - In other cases, businesses may prefer to have one major company pay for the operation of the system and act as a central governing authority
    - These differences in preference will influence the decision of whether to choose a public, private, or permissioned ledger
- decide whether or not value will be transferred within the system using cryptocurrency or cryptocurrency-like systems:
    - Adoption will depend on the industry context and each industry's willingness to employ currencies and systems that are currently still rather volatile

What type of smart contract and distributed ledger system is right for my business?

Once a problem has been clearly defined, and there is the possibility that smart contracts and distributed ledgers might be useful, businesses will need to design an approach that takes into account not only the features and limitations of the technology, but the features and limitations of the industry or system in which the problem exists. Using this understanding, they should make decisions by asking the right questions.

One of the main challenges that many companies are confronted with when figuring out how to implement distributed ledgers and smart contracts effectively is the tension between idealism and pragmatism – between what the technology could conceivably do given ideal conditions and what it can do now in the context of existing industries. Businesses must find an effective balance between these two competing approaches. If they do not, they will either limit the uptake of their proposed solution by being overly idealistic, or undermine the reason for using the technology in the first place by being overly pragmatic.

> *In order to build something that has an impact in the world it needs to be adopted. If nobody adopts your idea or product or system, it's not making an impact.*
>
> ***– Calogero Scibetta, Everledger***

Some within the blockchain community would argue that any compromise from the 'idealistic' fully trustless vision for distributed ledgers nullifies their potential, but there appear to be cases where a pragmatic approach will still result in implementations that help to underpin trust within interactions. Many of those we interviewed expressed the view that distributed ledgers and smart contracts can be used to build upon existing processes and relationships without needing to replace them entirely – for example, using permissioned ledgers to bring together existing ecosystems or recording transactions on a distributed ledger but using traditional bank transfers to process payments.

> *That's one reason why we actually don't take a role in verifying information because our intent isn't to try and make this market a perfect market, it's designed to make an improvement over the status quo.*
>
> ***– Jess Houlgrave, Codex Protocol***

Beyond just enabling an improvement over the status quo, some businesses believe that by adopting a more pragmatic approach they are helping their industry take steps that will eventually lead them to more 'ideal' or fully trustless approaches – albeit slowly and incrementally.

> *The reality is the world's supply chains are run by a small number of companies and it's not in their interest to facilitate a fully trustless system. But at the same time they're conflicted because they want to end things like modern slavery in supply chains and they don't want unsafe products – so they recognise that there's potential for that type of system. Part of what we do is slowly walking those big guys through the doors that could enable them to be part of a fully decentralised system – it's a journey.*
>
> ***– Jessi Baker, Provenance***

Smart contracts for business: selecting systems

The applicability of smart contracts depends firmly on the proposed use case and the ability of a business to define a problem, ask the right questions and remain pragmatic about a potential solution.

The vision offered by the advocates of smart contracts and distributed ledgers centres around the ability for businesses and others to interact more efficiently and effectively than they do today. At the heart of this promise is the belief that distributed ledgers can be used to underpin trust in these interactions in ways that other centralised technologies might not. On top of this, smart contracts potentially offer a means to automate business processes and interactions in a way that enables less reliance on other trust mechanisms.

In this report we have provided information we felt was important for businesses to know in order to make informed decisions about smart contracts. Our hope is that the information we have provided and the questions we have raised in Part 1 will enable businesses to ask some of the right questions when attempting to discern whether distributed ledgers and smart contracts could usefully be deployed within their businesses to solve real-world problems. In addition, we hope that the three challenges we have outlined in Part 2 will help businesses that have already decided to pursue smart contracts define the type of system that will most benefit their business and their clients.

As the technology develops and as the use cases for distributed ledgers and smart contracts crystallise and, potentially, become more stable, the questions businesses will need to ask themselves will change, as will the challenges that businesses will need to confront. For instance, if the infrastructure matures, businesses may be able to more readily identify real-world problems that are capable of being addressed by distributed ledger technologies and smart contracts. In a similar way, a more mature infrastructure might decrease the likelihood of crippling bugs and therefore convince businesses that storing smart contract code on the chain offers the best way forward. As a final example, if a stable coin emerges in the next few years, businesses – and, importantly, their clients – may feel more comfortable exchanging value using distributed ledgers and smart contracts. Only time will tell.

Regardless of how things develop over the coming years, the two related imperatives detailed above will remain: first, identify and define the problem to be solved before focusing on the technical solution; and second, remain pragmatic rather than idealistic about the capabilities of the technology and the context in which it is to be implemented. Businesses that follow these simple rules and prepare themselves to make decisions related to the three challenges identified above should be well prepared in their efforts to define whether, and how, distributed ledgers and smart contracts can help them tackle real-world business problems.

Appendix: Methodology

The goal of this report was to offer neutral, unbiased advice to businesses to help them better understand the potential of smart contracts and distributed ledgers. It was part of a three-year research project funded by Innovate UK, the UK's innovation agency.

To help us begin our research we commissioned Navin Ramachandran, UCL Centre for Blockchain Technologies and IOTA Foundation, and James Brogan, MD candidate at Albert Einstein College of Medicine and Research Fellow at UCL Centre for Blockchain Technologies, to conduct an overview of relevant literature related to distributed ledgers and smart contracts, focusing on smart contracts as a potential solution to business problems. In addition to detailing the key features and characteristics of smart contracts and distributed ledgers, their report explored the potential limitations of smart contracts, the requirements for a platform to be able to support smart contracts, and how to integrate smart contracts with other technologies and processes - among other topics.

Building upon this research, we conducted a landscape review of companies developing, implementing or using smart contracts. Businesses were targeted for interviews based on how they were attempting to implement smart contracts and the industries in which they were working.

We conducted 14 interviews during the course of our research, 12 of which were conducted in person while two took the form of written answers to questions exchanged over email.

We strove for a balance between people from academia, research institutions, startups and SMEs. Our interviewees were from companies headquartered in England, Northern Ireland, Singapore and the United States of America. As a reflection of the nature of the blockchain and smart contracts communities, many of the companies we spoke to were international in nature and scope with employees based in multiple countries around the world.