

Information Rules Smart Ledger Architectures & Distributed Permissions



November 2018



**Information Rules
Smart Ledger Architectures & Distributed Permissions**

November 2018

Maury Shenk

Managing Director, Lily Innovation

Professor Michael Mainelli

Executive Chairman, Z/Yen Group

Foreword

One needs permission from somebody or someone to do pretty much anything these days. Whether it be parking your car, gaining a license, accessing a stock exchange or, more recently and topically, retaining personal data from a client business card, somebody somewhere insists you provide evidence that you have permission. The information behind the permissions is maintained on a 'ledger' of some sort.

Many of these permissions are assumed or implicit, taken for granted or overlooked, particularly when considering how one buys and sells in a marketplace, be that a local vegetable stall, a commodities exchange, or an on-line goods platform.

In my time in investment banking, and then as an Alderman and later Sheriff of London, I have experienced countless examples of onerous documentation requirements and often wondered where this will end. 'Smart Ledgers' are multi-organisational databases with a super audit trail, and some embedded computer code. Smart Ledgers can provide flexible, technical architectures that can help us simplify the administration of these legally complex permissions.

I welcome this report that sets the scene, explains the technology, and attempts to suggest how Smart Ledger technology can enforce the 'information rules' our increasingly complex society demands.

A handwritten signature in black ink, appearing to read 'William Russell', with a stylized flourish at the end.

William Russell

Alderman for the Ward of Bread Street, City of London

Contents

Foreword	2
Preface.....	4
Introduction.....	7
1. The Co-Evolution Of Markets & Permissions.....	9
A. Physical Markets (Layer 1).....	11
B. Virtual Markets (Layer 2).....	14
C. Information Markets (Layer 3)	19
2. Permission Frameworks For Information Markets	25
A. Information & Communications Theory	25
B. Permission Frameworks	29
C. A Way Forward	37
3. Smart Ledgers: An Emerging Tool For Distributed Permissions	40
A. Technical Features	41
B. Legal Requirements.....	42
Conclusion	49
Appendix 1 - Distributed Ledgers & GDPR.....	50
Appendix 2 - Permission Framework Examples.....	56
Principal Authors	59
Acknowledgments	61

Preface

Smart Ledgers, mutual distributed ledgers with embedded computer code, are emerging as important tools for sharing information among people and organisations. This information takes the form of documentation of 'items', 'actions', and 'permissions'. Digital 'items' can be identity papers, health test information, pieces of music, company shares, or any other type of content or data. 'Actions' record what has been done, e.g. a purchase, a transfer, an addition, a deletion.

This paper explores 'permissions', what actions am I allowed to perform with items. At first this might seem simple. I (a patient) have a right to 'own my data'. In turn, doctors now have to ask permission to look at my data. You (a doctor) are allowed to look at my health data when I give you permission.

But what if I am ill? Do doctors still have to ask to look at my data? Of course. What if I am unconscious in an emergency ward? Well, of course, then my doctor has a right to look at my data without my direct consent at that time. What if my doctor is unavailable? Can another doctor look at my data? Can my doctor share my data with other doctors and specialists? Can the government look at my health data? For what purpose? Is my data destroyed when I die? Can I ask that my health data be handed to researchers upon my death? Can my offspring object to my health data being handed to researchers for genetic privacy reasons? Etc.

Similar questions arise in identity documentation. Can a bank hand my identity documentation to another bank? To a regulator? In what circumstances? Or music. Can I pass a piece of music on to another machine I own, to another person, can they in turn, how many times, who pays, when? Or internal business documents. Could an executive or senior manager in a financial services institution send important documents supporting crucial decisions, e.g. credit decisions, risk ratings, or rate setting, to an external 'internet of record' in a Smart Ledger for future retrieval after employment. In the event of regulatory investigation years after leaving the institution, he or she is personally liable. So, the terms of employment provide for being able to retrieve vital documents from the Smart Ledger after employment, when the executive and the firm's interests may diverge. If a regulatory investigation swings into action, a lawyer, holding an escrow decryption key, validates if it is a qualifying regulatory

investigation under the agreed terms with the institution and provides the key to the executive to use those documents in his or her defence.

Physical constraints on transporting paper documents or playing music helped to establish 'conventions' about the permissions we gave, and provided environments where use could be limited. Only so many people could see an old doctor's file on a patient. Only so many people could fit in a music hall. Technology changed all that - the photocopier, fax, computer, phonograph, radio, iPod. Digital items today come as discrete digital objects. Typically central third parties control databases that in turn provide permissions. These systems are typically quite crude. They require access to a central database, often provided by a password. The central database typically provides some basic access control, 'you can't go further without permission', and charging mechanism, 'viewing this report costs X'.

Because Smart Ledger technology more easily permits such permissions to be embedded as computer code, almost 'wrapping' each individual digital item, and for the use of permissions and items to be recorded as actions, new data 'swapping' markets and systems are being built today with Smart Ledger technology that permit great complexity in permissioning. Such systems can be very powerful compared with paper-based equivalents. Smart Ledger systems contrast with central third-party systems though in that each individual item can have a completely bespoke piece of code 'wrapping' around it.

Controlling the permissions surrounding digital items becomes more important. Such permissions rapidly become complex. This paper began when we asked ourselves how we could mentally encompass the world of permissions. There just seemed too many.

Ultimately, digital items, actions, and permissions find their way into computer code and data structures. This means that we need to instruct programmers about the types of permissions we want them to encode. Our core search became, "is there a way of stating permissions that we can express in a succinct and clear way to ensure good computer code and few misunderstandings?"

We wondered if this could take the form of a 'logic' or an 'algebra' that could express such permissions in a formal way, and perhaps permit formal proofs of

long chains of interacting permissions. This paper explores where we are in setting out those 'information rules'.

Maury Shenk & Michael Mainelli

Introduction

Markets depend upon *permissions*, which can be defined as:

“the action[s] of officially allowing someone to do a particular thing; consent or authorization.”¹

This definition is sufficiently broad to capture the wide array of permissions that are required in modern markets, so long as we read ‘officially’ to include non-governmental authority, *e.g.*, the authority of an employer or a business counterparty or a parent.²

For example, consider a London family that purchases a summer sailing holiday in Greece on a bareboat (*i.e.*, just the boat, no captain). Enjoying the benefits of the purchase requires a host of permissions. Even before setting sail, the family will need rights to fly on an airplane (represented by an e-ticket) and to sail the boat (both the contract with the boat hire company and likely some kind of sailing licence of a least one family member), passports to provide identification to the airline and the right to enter Greece, some kind of commercial credential or relationship (perhaps a credit card or bank account) allowing conversion of British pounds sterling into Euros, and many more. Once underway, the boat’s non-Greek-speaking passengers will need to negotiate berths in crowded Greek marinas, tables at only marginally less crowded portside restaurants, and the occasional souvenir purchases. For the latter at least, one can expect permission to be readily given!

Since the time of the first markets, before the beginning of recorded history, the complexity of such permissions has multiplied. This increase in complexity has accelerated in recent decades as information technologies have increased globalisation and de-materialisation of markets. In order to effectively manage such permissions in contemporary markets, we believe that a new framework for defining and managing distributed permissions is needed: what one might call a species of ‘information rules’.

¹ Oxford Living Dictionaries, <https://en.oxforddictionaries.com/definition/permission>

² ‘Official’ can be defined as “of or pertaining to an office, post, or place”. *Oxford English Dictionary* (1979: Oxford University Press).

The goal of this report is to consider how Smart Ledgers could implement such a permissions framework. We intentionally explore new technical approaches, rather than recommending an evolution of current frameworks and technology. We believe that current technical ‘architectures’, largely based around a central third party and its information technology, are giving way to more distributed architectures. These distributed architectures are based on ‘Smart Ledgers’, basically these are multi-organisational databases with a super audit trail, typically containing some embedded computer code.

First, we set out a framework for analysing market permissions (as well as the related concepts of obligations and prohibitions), aiming to show that there is need for robust thinking about how permissions should work in information-driven markets. Second, we consider possible approaches to a permission framework for such markets and propose a way forward. Third, we explore how Smart Ledgers can be an important technical approach (although not an exclusive one) for implementing such a permission framework, and we comment on associated technical and legal challenges (*e.g.*, issues that arise under the new EU General Data Protection Regulation).

1. The Co-Evolution Of Markets & Permissions

Humans have exchanged goods since before the beginning of recorded history. Markets, which are an important innovation of human exchange, are also very old. What distinguishes markets from simple exchange is a degree of organisation. We turn again to the dictionary for a definition of ‘market’:

“a regular gathering of people for the purchase and sale of provisions, livestock, and other commodities”

or

“an area or arena in which commercial dealings are conducted”.³

Although it is common in contemporary discourse to speak of ‘the market’ – as if it were a single thing – in fact markets have a wide variety of forms, differing by purpose, geography and other factors. There is little superficial resemblance among a cattle auction in Argentina, children swapping trading cards on the playground, and the competition of independent taxi drivers for rides on Uber and Lyft.

Yet there are common features of markets. It is a central thesis of this paper that, on the winding path from prehistoric exchange to the present day, markets have developed three ‘layers’: (1) a *physical layer* supporting actual exchange of goods and services (sometimes over great distances), (2) a *virtual or communications layer* allowing negotiation of transactions to be separated in time and space from physical exchange, and (3) an *information layer* in which metadata regarding exchange becomes valuable separate from the exchange itself.

These three layers developed (mostly⁴) sequentially. The physical layer is the oldest, dating to the beginning of markets; the virtual layer has been important for just over a thousand years; and the information layer is only a few decades old. It is this relative novelty of information markets that gives rise to the issues in this report, because the rules of the game for information markets continue to be written as those markets evolve rapidly.

³ Oxford Living Dictionaries, <https://en.oxforddictionaries.com/definition/market>

⁴ Critical readers will likely find analogues of what we call ‘virtual markets’ or ‘information markets’ even in antiquity.

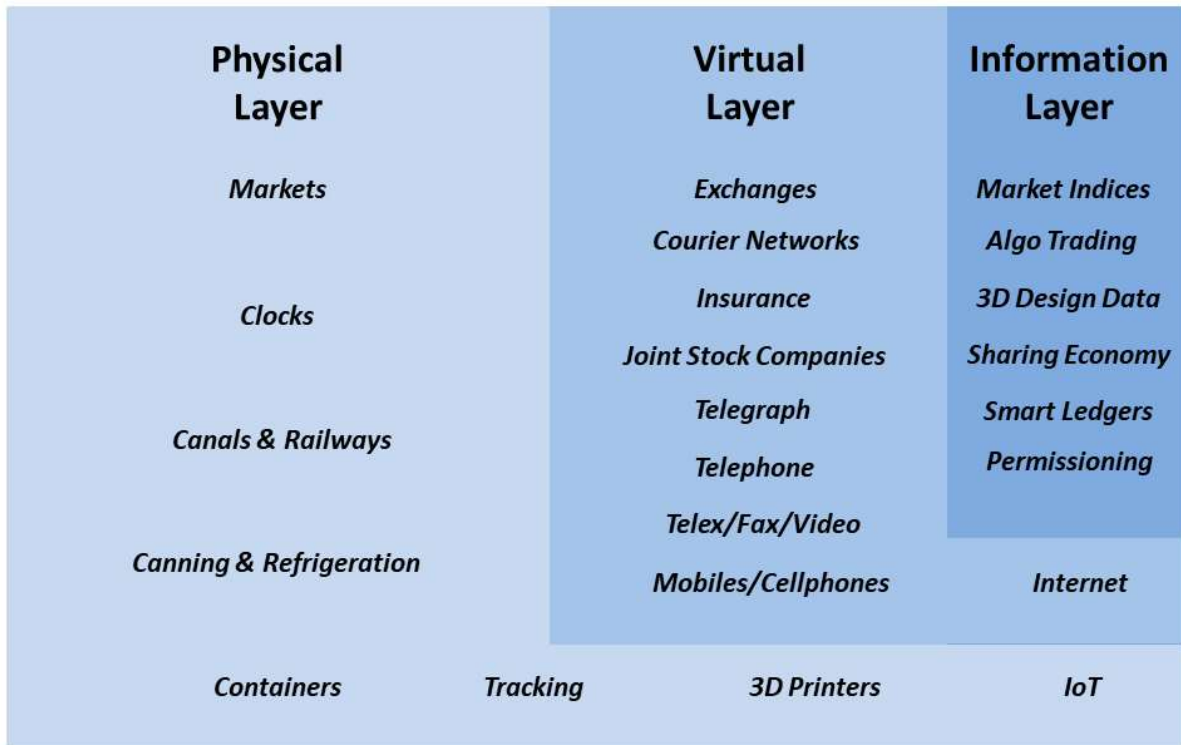


Figure 1 - Physical, Virtual, and Information Markets

At each layer, permissions are a crucial aspect of market rules. However, the nature of permissions differs significantly among the layers, for the simple reason that access to a physical market square (and purchases there) involves a very different set of issues from transactions by telecommunication (at the virtual layer) or exchange of information about market function (at the information layer).

Before turning to permissions, we briefly consider the concept of ‘assets’ – *i.e.*, the stuff that is traded on markets. Turning once more to the dictionary, ‘asset’ can be defined as “a useful or valuable thing or person”.⁵ But this straightforward definition is a bit rudimentary for our analysis of permissions. One of the authors identified seven criteria that can be used to evaluate an asset from the perspective of an auditor (spelling out the fishy acronym ‘COD-VERB’):

- accurate understanding of the Cost of the asset;
- confirmation of Ownership of the asset;
- some Disclosure of the importance of the asset;
- ability to confirm the Value of the asset;

⁵ Oxford Living Dictionaries, <https://en.oxforddictionaries.com/definition/asset>

- evidence of the Existence of the asset;
- clear lines of Responsibility for the asset;
- measurable Benefit from the asset.⁶

A detailed analysis of assets based on the COD-VERB framework is beyond the scope of this report, but the essence of the framework provides a suitably granular background for considering assets in the context of permissions. Two pairs of related criteria – ownership and responsibility, and value and benefit – are central to our analysis below of assets and permissions at each of the three market layers. That is, we focus on which actor has ownership and responsibility for assets in each market, and how they derive value and benefit from those assets based upon appropriate permissions.⁷

A. Physical Markets (Layer 1)

Markets for the trading of physical goods have existed since before the beginning of recorded history. For example, there is evidence of markets in southern Europe in the Bronze Age (roughly 2500 – 1000 BCE) and in the Middle East as early as 3000 BCE.⁸

The reason for early development of markets is obvious. Where goods are traded among individuals, it is efficient for trade to happen at a central location, rather than buyers and sellers needing to seek out their counterparties across the landscape (at a time before convenient communication or directories). Eventually, physical markets developed to allow trade over great distances, such as the Silk Road (the ancient network of trade routes connecting the Mediterranean region with China). Here is a description of physical market transportation within China at the glorious height of the Tang Dynasty in the 8th century:

⁶ Michael Mainelli & Ian Harris, “All Important Information”, Z/Yen Group publication (2002), <http://www.zyen.com/37-publications/professional-articles/160-all-important-information-intelligence-articles-2002.html>

⁷ The remaining three criteria of existence, cost and disclosure also make appearances below, including in the context of permissioning algebra that is discussed in section 2 below.

⁸ “Marketplace: Markets in prehistory”, *Wikipedia*, https://en.wikipedia.org/wiki/Marketplace#Markets_in_prehistory

“Mules and horses were available to travellers on ... secure roads, and an intricate system of canals devised to provide water transport for tax silks from the mouth of the Yangtze River to the capital [Chang’an (modern day Xian)] was now so improved that it could also be used to bring luxury goods from foreign countries.”⁹

In these layer 1 (physical) markets, the main relevant assets are the goods being traded, and those goods have value primarily because of their utility to their owners. Services can also be traded in layer 1 markets (*e.g.*, a haircut, or the offerings of the oldest profession), although the treatment of services as distinct from goods is a comparatively recent phenomenon – as recently as the early 20th century, leading economist Alfred Marshall wrote that services are simply “immaterial products”.¹⁰

The key permission in layer 1 physical markets for goods is *ownership*. This can be defined as an inherent characteristic of the asset itself (as in the COD-VERB model), reflected in the truism that “possession is nine-tenths of the law”. However, ownership can also involve elements of official permission, including deeds of title (*e.g.*, in the case of an automobile) and judicial determination of ownership in the event of a dispute. Furthermore, permissions for physical markets can extend well beyond ownership. We consider two examples: market cities of Middle Ages Europe, and the modern convenience shop.

Market Cities

The Venerable Bede, a monk and historian living and working in the late 7th and early 8th century in a monastery in England’s Northumbria, wrote that London of the early 7th century was “the mart of many nations resorting to it by land and and sea”.¹¹ Though written records are fairly limited, there is other extensive evidence of the development of London as a trading port during the following centuries.

⁹ Edward H. Schafer, *The Golden Peaches of Samarkand*, p. 8 (1963: University of California Press).

¹⁰ Alfred Marshall, *Principles of Economics* (8th ed. 1920: Routledge).

¹¹ Gustav Milne, *The Port of Medieval London*, p. 30 (2003: Tempus Publishing). Though Bede’s travels from his monastery were limited, he managed to synthesise work of historians and travellers to write intelligently about the history and actuality of England, and rather less accurately about affairs further afield such as in Spain and the Near East. Katharine Scarfe Beckett, *Anglo-Saxon Perceptions of the Islamic World*, pp.18-21 (1999: Cambridge University Press).

Far better records of market development are available about later markets, such as the ‘Lion City’ of Venice during the late Middle Ages, as historian Peter Ackroyd describes:

“Venice possessed no natural resources, and so it relied upon manufacture; the only way of maintaining its supremacy was in the creation of more various and more rarefied items. Luxury was prodigality, whether in spices or perfumes or dye-stuffs or ornaments of gold and rock crystal. Venice traded in them all. It made the glass and the silk and the soap. It manufactured the marzipan as well as the wax. Venice was a centre of silk manufacture, while the neighbouring island of Burano was the home of lace-making and Murano of mirrors and glass.”¹²

The main feature of the markets of London and Venice was simply their wide variety of goods. And, as is the general case for layer 1 markets, the key permission for trade in market cities was ownership of goods. However, the right and ability of merchants to trade also required permission from the local sovereign. In the case of Venice, this resulted in regular military conflicts for control of trading territory and trade routes with other city states – particularly Genoa, the great rival of Venice – and great walls were built to defend the city.¹³

Convenience Shops

The modern convenience shop (which is rapidly ceding commercial territory to Amazon.com, Alibaba.com, eBay.com, and other online competitors) has significant similarities to a medieval market, offering the possibility to buy a wide variety of goods (and some services) in a single location. Most of these products do not require permission (beyond ownership) to be sold, with some exceptions. For example, in most countries sales of alcohol require a licence, and sales of both alcohol and cigarettes typically require proof that the buyer is above a certain age (18 years in the UK). Furthermore, the owner of the convenience shop likely requires a rental contract for the physical space, and perhaps a franchise agreement if part of a larger chain. Permission is also required to support the constantly-evolving electronic payment methods now available to

¹² Peter Ackroyd, *Venice: Pure City*, pp. 128-29 (2010: Vintage Books).

¹³ Ackroyd (note 12 above), pp. 204-22.

consumers – which leads us to the virtual layer of markets, and commerce at a distance.

B. Virtual Markets (Layer 2)

The key innovation offered by virtual markets was to provide means for merchants to agree exchanges of goods at a place separated in space or time from the physical exchange of goods. In Europe, this seems to have occurred in the late Middle Ages. The Flemish ‘beurs’ and French ‘bourse’, meaning ‘exchange’, appear to have originated from a hanging sign outside an inn in the Belgian city of Bruges in the 13th century:

“One of the most important families of innkeepers [in Bruges] was the Van der Buerse family. For five generations they ran the ‘Ter Buerse’ inn. The oldest records relating to the family date from the 13th century, and it is an established fact that the Ter Buerse inn itself was already operating in 1285. It was run by Robert Van der Buerse, who was also the owner. During the 14th century the square in front of the Ter Buerse inn developed into the leading commercial and financial centre of the city. By 1340, Pegolotti’s guide to commerce was already comparing the Bruges brokerage rates and exchange rates with trade in England and Italy. Brokers also met in the square at set times, and because there were not yet any official stock market journals, they gathered a whole range of information from their guests, who travelled and corresponded with them about the local economic situation and the state of the foreign markets. In 1370, exchange rates for various cities were announced at regular times in Bruges. By around 1400 a continuous, organised money market had developed, with set times for announcing the exchange rates of the leading commercial and banking centres of Europe, such as Barcelona, Venice, London and Paris.”¹⁴

¹⁴ National Bank of Belgium, *The Stock Market: From The ‘Ter Buerse’ Inn To Wall Street*, 2010 - <https://www.nbbmuseum.be/en/2010/01/stockmarket.htm>

Bourses spread throughout the low countries. In London, Sir Thomas Gresham (1519-1579) imported from Antwerp the idea of a 'bourse' or 'exchange', a permanent market for intangible items such as ship voyages and insurance, as well as trading bills of exchange. Incorporated above the 1571 Royal Exchange were 150 small shops, called The Pawn, London's first shopping centre.^{15,16} Arguably, both the London Stock Exchange and Lloyd's of London sprang directly from this venture.

In layer 2 virtual markets, the main relevant assets are such trading venues, and means of exchange and communication between them. These assets have value because of their ability to improve the efficiency of underlying layer 1 markets. The key permission in layer 2 markets is *access* to trading facilities and platforms. We again use an historical example (the Rothschild banking system) and a modern one (high-frequency trading) to illustrate.

Rothschild Banking Network

Echoing the origins of the word 'bourse' in Bruges, the Rothschild family name derives from the German '*zum rothen Schild*' ('at the red shield'), which referred to a sign on the Rothschild family house in Frankfurt. It was from Frankfurt that Mayer Amschel Rothschild established the basis for the Rothschild banking dynasty by sending his five sons to the five main European financial centres (Frankfurt, Vienna, London, Naples and Paris) to build a banking network.¹⁷

The Rothschild banking network, and the multinational transactions that it supported, required effective communication among the Rothschild banks and other locations. The Rothschilds had the most effective communications in Europe, originally built to smuggle gold bullion and later enhanced to support a wide variety of financial trading operations:

“[T]he success of arbitrage and forward exchange operations hinged on rapid communication. As far as possible, the brothers sought to keep one

¹⁵ Alfred Edward Woodley Mason, "The Royal Exchange : A Note On The Occasion Of The Bicentenary Of The Royal Exchange Assurance" (1920, Royal Exchange Assurance) - <https://archive.org/details/royalexchangenot00masoiala>

¹⁶ David Kynaston, *City of London: The History*, p. 9 (2012: Vintage Books).

¹⁷ "Rothschild family", *Wikipedia*, https://en.wikipedia.org/wiki/Rothschild_family. See also generally Niall Ferguson, *The House of Rothschild* (1999, Penguin).

another abreast of news which might affect the exchange markets: the impending payment of a new subsidy, the likelihood of further military action, the imminence of the peace treaty being signed. And ... they were already able [during the Napoleonic Wars] to transmit such information through their own couriers considerably faster than was possible through official channels or the regular post. Yet the time-lags could still be substantial and Nathan [Rothschild, based in London,] was constantly being urged to speed up the system.”¹⁸

This communications system is a canonical example of market layer 2 – enabling reliable transactions between economic actors in different locations, without simultaneous exchange (although with time lags reduced to the extent possible). The relationship between the layer 2 communication network and the underlying layer 1 banking and financial transactions is crucial. Specifically, the transactions that the Rothschilds were able to facilitate using their communications network shared many characteristics with transactions negotiated face-to-face, but crucially *the communications network enabled transactions that would have been impossible in its absence*. By adding layer 2 to their banking operations, the Rothschilds were able to arrange transactions of a scope and scale never before seen – notably playing a lead role in financing the British war effort in the Napoleonic Wars as well as the independence of Brazil from Portugal, both in the early 19th century¹⁹ – and to become tremendously wealthy in the process.

The layer 2 assets underlying the Rothschilds’ success were thus relatively straightforward: their trading facilities around Europe and the communications channels between them. Less straightforward were the layer 1 assets – gold bullion, company shares and many other assets – that were traded by the Rothschild family. It required a tremendous combination of intelligence, daring and good luck for the family to accumulate the underlying assets that it traded. The permissions that the Rothschilds required at layer 2 were also linked to the layer 1 assets. Specifically, political and commercial connections provided permissions that the Rothschilds needed to engage in transactions at the highest levels of European commerce, even though the family was a Jewish one that initially struggled to be accepted by the European aristocratic order. Tying

¹⁸ Niall Ferguson, *The House of Rothschild* (1999, Penguin), p. 94.

¹⁹ Niall Ferguson, *The House of Rothschild* (1999, Penguin), pp. 83-110, 132-33.

together assets and permissions, historian Niall Ferguson attributes the emergence of Nathan Rothschild as the key financier for the British war effort against Napoleon to three factors:

1. most importantly, the communications and logistics infrastructure to get money to Wellington and his armies (*i.e.*, the layer 2 assets);
2. connections to (and the favour of) John Charles Herries, Commissary-in-Chief of the British Army (providing the needed permissions);
3. the good fortune to face limited competition after the financial crisis of 1810 in England, and the collapse of the Amsterdam market caused by Napoleon's annexation of the Netherlands.²⁰

High-Frequency Trading

In the two centuries since the Rothschilds attained prominence, the importance of rapid communications has remained a constant in financial markets, which are now almost exclusively layer 2 markets (practically no one buys stocks or bonds face to face any more). Speeds have increased by many orders of magnitude. High-frequency trading ("HFT") firms, who seek to exploit small, temporary price differences for the same (or linked) securities on different markets, account for much of the trading on modern securities markets. Where the Rothschilds were seeking time advantages of days or hours, HFT firms trade on time advantages of milliseconds (thousandths of a second) for exchanges separated by significant distances, or microseconds (millionths of a second).²¹ To some degree, this advantage is due to a change around 2000 from markets that tried to 'match' orders of various sizes (*e.g.*, Stock Exchange Electronic Trading Services (SETS, yes there should probably have been two 'E's) in the London Stock Exchange) moving to 'first in, first matched, first out' trading. This is a subject of considerable discussion, not least Long Finance's proposal for a global 'fair exchange' where trades within a sizable time block, perhaps 10 seconds or more, could be matched from around the world fairly. Investors Exchange (IEX), founded in 2012 in the US, has taken many 'fair play' ideas forward and has grown its market share to over 2% already.

²⁰ Niall Ferguson, *The House of Rothschild* (1999, Penguin), pp. 85-87.

²¹ See generally Michael Lewis, *Flash Boys: Cracking the Money Code* (2015: Penguin).

The main constraint on speed for HFT firms is the speed of light in optical fibre, *i.e.*, about 125 miles per millisecond, possibly somewhat faster.²² Maximising speed means minimising distance. This has driven a race by HFT firms to locate their facilities and computers close to stock exchanges, and to have access to the straightest communication routes between exchanges. For example, a company called Spread Networks spent hundreds of millions of dollars about a decade ago to build as straight as possible a fibre optic line (including tunnelling through mountains) between New York and Chicago.²³ Because portfolio managers need to get good prices and cannot sell large bundles of shares safely, they need to break up their orders and often ‘buy’ rather than ‘sell’ to hide their overall strategy. This results in a lot more trading. It is difficult to see how this race is adding to value, as several studies show that HFT leads to overall increased costs to manage a portfolio, though individual trades are cheaper and have more liquidity, except during ‘flash crashes’.²⁴ There are also significant energy processing cost implications for the environment.²⁵

Thus, for HFT firms, as for the Rothschild family, the key layer 2 assets are trading facilities and the communications links between them. However, crucial permissions have shifted. Although some players in modern financial markets (such as the big investment banks) still depend heavily upon access to customers, the advantages of HFT firms derive largely from rights to locate computer equipment at desired *locations* that are closer to exchanges than their competitors’ computers.

At the same time, regulatory permissions have become much more complex over the centuries. The financial markets in which HFT firms operate are subject to detailed regulation, by bodies such as the Financial Conduct Authority in the

²² See Brian Quigley, “Speed of Light in Fiber – The First Building Block of a Low-Latency Trading Infrastructure”, ADVA Optical Networking blog (April 7, 2011), <https://blog.advaoptical.com/en/speed-light-fiber-first-building-block-low-latency-trading-infrastructure>.

²³ Michael Lewis, *Flash Boys: Cracking the Money Code* (2015: Penguin), pp. 8-22.

²⁴ “How High-Frequency Trading Hit A Speed Bump”, FT (1 January 2018) - <https://www.ft.com/content/d81f96ea-d43c-11e7-a303-9060cb1e5f44>

²⁵ Michael Mainelli, “Green Data: Impact Versus Low Latency?”, *Inside Market Data*, Volume 24, Number 32, Incisive Media Limited (11 May 2009), page 9 - <https://www.zyen.com/publications/professional-articles/green-data-impact-versus-low-latency/>

UK and the Securities and Exchange Commission in the United States. However, HFT firms are relatively lightly regulated within these markets – typically requiring only authorisation as a broker-dealer and some other basic formalities in order to trade, and often benefitting from their ability to game the regulatory system.²⁶

C. Information Markets (Layer 3)

But high-frequency trading is about more than just faster communications. Equally if not more important are the software algorithms that make trading decisions based upon the just-in-time market data provided by the underlying super-fast communications networks. It is such technology for the manipulation of information about market exchange – separate from the exchange itself – that is the basis of layer 3 information markets. To illustrate what is different about layer 3, consider that a high-frequency trading algorithm might evaluate price movements of many thousands of assets or securities before making buy or sell decisions about a tiny fraction of them. The underlying trade in the selected assets might resemble a trade of many decades ago, but the overall trading strategy was simply impossible in the absence of computer technology. Before layer 3 is consigned to financial markets such as share trading or online gambling, remember that online advertising markets have analogously almost all the features of financial markets. Thus, the introduction of layer 3 technology has been revolutionary much in the way that the Rothschilds' introduction of layer 2 technology to banking was revolutionary.

Information markets are largely driven by computer technology, so began to flourish in the latter part of the 20th century. However, information markets began to emerge earlier. For example, Charles Dow created the Dow Jones Industrial Average, which provides metadata on the United States stock market, in 1896 – 12 years after Dow created his first a market index, focused on transportation.²⁷ Dow's indices provide layer 3 metadata about transactions on layer 2 stock exchanges. But notwithstanding such early examples, the true

²⁶ Michael Lewis, *Flash Boys: Cracking the Money Code* (2015: Penguin), pp. 96-99 (describing how HFT firms have taken advantage of the pricing information created by United States Securities and Exchange Commission Reg NMS, which requires brokers to seek the 'best price' for their clients).

²⁷ "Dow Jones Industrial Average", *Wikipedia*,
https://en.wikipedia.org/wiki/Dow_Jones_Industrial_Average.

potential of layer 3 markets is only now being realised with the development and diversification of information technology.

In layer 3 information markets, the main relevant assets are software, algorithms and computers, with software and algorithms accounting for most value creation (computers are now commoditised). As Marc Andreessen, leader of the team that developed the Mosaic, and later Netscape, internet browser, famously wrote in 2011, “software is eating the world”.²⁸ Just as layer 2 assets have value because of their ability to improve the efficiency of layer 1 markets, layer 2 assets have value because of their ability to improve the efficiency of layer 1 *and* 2 markets. The key permissions in layer 3 markets involve access to information and rights to use it, *i.e.*, intellectual property, defined broadly. We illustrate the blossoming of information markets with two examples of contemporary Internet business models: on-demand automobile transportation (Uber, Lyft, and their competitors) and intellectual property acquisition and licensing (Intellectual Ventures).

Transportation Exchange - Uber And Lyft

Uber, Lyft, and their competitors are revolutionising the global transportation market²⁹ – and devastating the market for traditional taxicab services – in a way that would have been impossible before the advent of the smartphone. But it is not only the connectivity features of the smartphone that enable this business model. Uber and Lyft do not operate at layer 1 (where individual drivers provide transportation services), and much of their operations at layer 2 (*i.e.*, connecting drivers with passengers) are nothing particularly new – telephone and online taxi booking services have existed for decades. Further, patent filings, such as those by Eric Masaba for a “Taxi Dispatch System” from 2006 show that the ideas are older than Uber and Lyft.³⁰

Two main things make the services implemented by Uber and Lyft revolutionary. First, the essence of their service is managing information about the locations of drivers and passengers, and matching them efficiently, including by using

²⁸ Marc Andreessen, “Why Software Is Eating the World”, *The Wall Street Journal* (Aug. 11, 2011).

²⁹ In some countries, local competitors are stronger than Uber, such as DiDi in China, Ola in India and Careem in the Middle East.

³⁰ <https://patents.google.com/patent/US20080015923#patentCitations>

dynamic price increases to increase driver supply (and reduce passenger demand) in areas where supply and demand are unmatched. This information service lies at market layer 3 – that is, both passengers and drivers are willing to cede part of their fare to Uber or Lyft in exchange for providing the information that makes the service possible. Second, the companies have innovated at layer 2 by requiring that fares be paid automatically to the passenger's credit card, making the entire service almost magically convenient, and facilitating allocation of payment between the application provider and the driver.

The key assets of on-demand transportation companies are fairly obvious: their algorithms and the software that implements them, and their networks of passengers and drivers. But permissions are more complex for these businesses, including:

- authorisation to operate a taxi service – Lyft was the first of the on-demand transportation companies to take the bold step of effectively providing taxicab service (albeit at layer 3!) without complying with taxi licensing laws. Even Uber's famously bold former CEO Travis Kalanick was reportedly reluctant to take the same risk, but promptly followed suit once Lyft had done so.³¹ Licensing battles – such as Uber's protracted effort to retain the right to operate in London³² and outright bans of its service in Denmark, Hungary and Bulgaria³³ – remain a major ongoing issue for these companies.
- ability to treat drivers as contractors not employees – Another major legal issue that Uber and Lyft face across multiple jurisdictions is whether they can continue to treat drivers as independent contractors, rather than higher-cost employees. For example, decisions in November 2017 by the UK Employment Appeals Tribunal and in May 2018 by the California Supreme Court (involving a different company) are among those that indicating that Uber and Lyft drivers are employees.³⁴

³¹ See Tim O'Reilly, *What's the Future and Why It's Up to Us* (2017: Random House Business Books), pp. 54-55.

³² Gian Volpicelli, "Uber's London licence has been approved – but there's a big catch", *Wired* (June 26, 2018), <https://www.wired.co.uk/article/uber-london-licence-tfl-verdict>

³³ Greg Dickinson, "How the world is going to war with Uber", *The Telegraph* (26 June 2018), <https://www.telegraph.co.uk/travel/news/where-is-uber-banned/>

³⁴ See Sarah O'Connon & Aliya Ram, "Uber loses appeal in UK employment case", *Financial Times* (Nov. 10, 2017), <https://www.ft.com/content/84de88bc-c5ee-11e7-a1d2-6786f39ef675>; "Uber and

- authorisation of drivers – In order to qualify as an Uber or Lyft driver, a car owner must provide a limited amount of documentation to Uber and complete screening and background checks.³⁵ An active driver receives passenger ratings and may be terminated if his/her average rating is too low (reportedly below about 4.6 stars for Uber).³⁶
- authorisation of passengers – In order to use Uber or Lyft as a passenger, an individual need only download the mobile application, and provide basic identifying information and payment information.³⁷
- control over data – Municipalities have encouraged Uber to make data on usage of its services publicly available, and Uber has done so for selected services.³⁸

A number of other ‘sharing economy’ or ‘sharing asset’ examples could be similarly highlighted from property (Airbnb or Homestay) to yachts (Boatsetter or Click&Boat).

Intellectual Property Exchanges - Intellectual Ventures

Another, slightly older layer 3 business model was pioneered in 2000 by Intellectual Ventures, which is seeking to build an “invention marketplace”

Lyft drivers could get employment status under California court ruling”, *The Verge* (1 May 2018), <https://www.theverge.com/2018/5/1/17308178/uber-lyft-drivers-california-court-classification-dynamex>

³⁵ Driver requirements: How to drive with Uber, <https://www.uber.com/drive/requirements/>; Lyft, Drive toward what matters, <https://www.lyft.com/drive-with-lyft>

³⁶ See Uber, Star ratings: What to know as a driver partner, <https://www.uber.com/drive/resources/how-ratings-work/>; James Cook, “Uber’s internal charts show how its driver-rating system actually works”, *Business Insider* (11 February 2015), <https://www.businessinsider.com/leaked-charts-show-how-ubers-driver-rating-system-works-2015-2>

³⁷ How do I create a Uber account?, <https://help.uber.com/h/fdfc0273-f67e-4545-9885-f89d0ca0aacf>; Riding with Lyft, <https://help.lyft.com/hc/en-us/categories/115002006488-Riding-with-Lyft>.

³⁸ Uber Movement, [https://movement.uber.com/cities?lang=en-United States](https://movement.uber.com/cities?lang=en-United%20States) (Uber traffic data portal); see Mariella Moon, “Uber Movement’s traffic data is now available to the public”, *Engadget* (Aug. 31, 2017) <https://www.engadget.com/2017/08/31/uber-movement-traffic-data-website-launch/>

centred on patent rights.³⁹ The company's founder Nathan Myhrvold, former Chief Technology Officer at Microsoft, described the company's business model in a 2010 article in *Harvard Business Review*:

“My company, Intellectual Ventures, is misunderstood. We have been reviled as a patent troll—a renegade outfit that buys up patents and then uses them to hold up innocent companies. What we're really trying to do is create a capital market for inventions akin to the venture capital market that supports start-ups and the private equity market that revitalizes inefficient companies. Our goal is to make applied research a profitable activity that attracts vastly more private investment than it does today so that the number of inventions generated soars.”⁴⁰

Intellectual Ventures and Myhrvold explicitly use 'market' language, and in our framework for markets their service is at layer 3. The information of value that is traded involves patent rights, which provide a legally-created monopoly (usually for 20 years) to practice an invention. In this context, layer 1 is the delivery of an underlying product or service, layer 2 is the virtual trading of the product or service, and layer 3 is the trading of intellectual property rights (usually patent rights) that authorises trading at layer 1 and/or layer 2.

The assets required for this business are primarily the intellectual property rights themselves, plus the software, facilities and human expertise that constitute Intellectual Ventures' 'invention marketplace'. The relevant permissions are tightly linked to the assets – *i.e.*, ownership of intellectual property rights, and the sub-divided permissions that Intellectual Ventures can grant through licenses to use those rights. In the language of COD-VERB that began this section, ownership and responsibility for intellectual property rights is tightly linked to value and benefit.

It can be easily observed that the assets and permissions in the two above examples at layer 1 (market cities and convenience stores) share important common features despite their separation in time, and the same is true for the

³⁹ Intellectual Ventures, About Us, <http://www.intellectualventures.com/about/invention-marketplace>.

⁴⁰ Nathan Myhrvold, "The Big Idea: Funding Eureka!", *Harvard Business Review* (March 2010), <https://hbr.org/2010/03/the-big-idea-funding-eureka>

two examples at layer 2 (Rothschild banking network and high-frequency trading). Although layer 1 and 2 markets are not always so similar, they tend to share more common features than do layer 3 markets. In the latter context, our two examples above (transportation exchanges and intellectual property exchanges) have very little in common even though they are both contemporary business models.

It is the task of the remainder of this report to consider how to develop a coherent framework for the wide variety of permissions in such information-based layer 3 markets. Section 2 below articulates general principles for such permission frameworks, and section 3 focuses on implementation of such frameworks using Smart Ledgers.⁴¹

⁴¹ We have intentionally not used Smart Ledger business models in our examples above, because a main goal of this paper is to show how Smart Ledgers can be used to implement permissions across a wide range of business models, including those discussed above.

2. Permission Frameworks For Information Markets

Permission frameworks for physical and virtual markets have been centuries (or millennia) in development. Although these frameworks continue to evolve slowly, they are generally well understood. For example, the concept of ownership is very old, and rules of contract law have evolved gradually over the centuries. Although financial market rules are more recent, even the modern high-frequency trading market described above operates within the framework of financial markets regulation largely developed in the early 20th century, gradually evolving since then.

By contrast, information markets are changing fast, and we are still in the early days of development of their permission frameworks. Taking the examples above, Uber, Lyft, and their competitors are able to contend that they are outside traditional regulatory frameworks for taxicab services because of the novelty of their business models; and the important reputation markets associated with their services are subject to few clear rules at all. Intellectual Ventures operates within the established regulatory framework for patents, but the interactions and permissions of their ‘invention marketplace’ are being established as that marketplace develops.

Furthermore, information markets run at rapid speed on information technology, too fast for humans to make individual decisions about permissions (this is also the case for modern layer 2 virtual markets). Any clear permission framework operating at computer speed (including on a Smart Ledger) must be unambiguous in a way that allows it to be computable.

These factors of rapid change and computational speed indicate a need to rethink permission frameworks in a way that is suitable for layer 3 information markets, while (ideally) at the same time covering traditional markets and the layer 1 and 2 aspects of information markets.

A. Information & Communications Theory

Before exploring potential permission frameworks that are suitable for information markets, we first should consider what is meant by ‘information’. There are many examples of information assets in current markets, such as:

- personal data – which is the fuel of the attention-driven business of companies like Facebook and Google;
- patents and other intellectual property – which is central to practically any technology business model, and is traded by companies like Intellectual Ventures;
- training data for artificial intelligence ('AI') – which is one of the two key drivers, along with expanded computing capacity, of the current explosive success of AI business models.

However, we need more than an enumeration of such examples to build a permission framework. What is essential is a general approach to 'information' that can handle the full spectrum of information assets that will be relevant to rapidly-evolving information markets.

As a starting point, we return to the dictionary for definitions of 'information':

“Facts provided or learned about something or someone.”

or

“What is conveyed or represented by a particular arrangement or sequence of things.”⁴²

These definitions, while common-sensical, are insufficient for our purposes. To build a somewhat more complete foundation, we take a very shallow dive into the disciplines of information theory, communications theory and the philosophy of information (a full summary of these disciplines is well beyond the scope of this report).

Information theory, which was first proposed by Claude Shannon in 1948, concerns the technical aspects of communication of messages:

“The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have *meaning*; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is

⁴² Oxford Living Dictionaries, <https://en.oxforddictionaries.com/definition/information>

one *selected from a set* of possible messages. The system must be designed to operate for each possible selection, not just the one which will actually be chosen since this is unknown at the time of design.”⁴³

Our primary concern in designing a permission framework for layer 3 markets is in fact *meaning*, which Shannon identifies as irrelevant to information theory. However, his distinction between *message* and *meaning* is nevertheless important to our inquiry. Specifically, we must ensure that any permission framework is sufficiently specific and flexible to convey precisely the content of any permission that we might wish to articulate, irrespective of its meaning.

Turning to meaning itself, we move from information theory to communication theory, and specifically to the contributions of the Toronto School, which included Harold Innis and Marshall McLuhan. Innis and McLuhan recognised that different ‘media’ can have very different properties. Innis focused on communications media, and postulated distinctions between ‘time-biased’ and ‘space-biased’ media:

“The concepts of time and space reflect the significance of media to civilization. Media that emphasize time are those durable in character such as parchment, clay and stone. The heavy materials are suited to the development of architecture and sculpture. Media that emphasize space are apt to be less durable and light in character such as papyrus and paper. The latter are suited to wide areas in administration and trade.”⁴⁴

McLuhan went a step further in defining a medium to include “any extension of ourselves”:

“In a culture like ours, long accustomed to splitting and dividing all things as a means of control, it is sometimes a bit of a shock to be reminded that, in operational and practical fact, *the medium is the message*. This is merely to say that the personal and social consequences of any medium – that is, of *any extension of ourselves* – result from the new scale that is introduced

⁴³ C. E. Shannon, “A Mathematical Theory of Communication”, *Bell System Technical Journal*, Vol. 27, p. 1 of reprint (July, October 1948) (original emphasis).

⁴⁴ Harold Innis, *Empire and Communications*, p. 27 (1950: Oxford University Press).

into our affairs by each extension of ourselves, or by any new technology.”⁴⁵

The central import of these views of Innis and McLuhan for our exploration of permission frameworks is that the *how* of human (or machine) interactions (*e.g.*, the communications medium used) has a significant effect on the practical consequences of those interactions.

This recognition that use and communication information has complex practical dimensions has been explored in significant further depth in the past two decades in the field of philosophy of information, founded by Oxford philosopher and ethicist Luciano Floridi:

“The philosophy of information investigates the conceptual nature and basic principles of information, including its ethical consequences. It analyses problems in order to design solutions. It is a thriving new area of research, at the crossroads of epistemology, metaphysics, logic, philosophy of science, semantics, and ethics.”⁴⁶

Floridi intended this discipline from the outset to be “capable of dealing with contemporary and lively issues about which we really care; and less prone to metaphysical armchair speculations and idiosyncratic intuitions ... a constructive philosophy which would provide answers, not just analyses”.⁴⁷ In this pursuit of applicability, he identifies various categories of information: mathematical information, semantic information, physical information, biological information, economic information, and sub-divisions of these broad categories.⁴⁸ Taking a somewhat different slant on similar issues, one of the authors of this paper has postulated that information can play various different roles, including those of memory, communication, intellectual property, market enabler and context.⁴⁹

⁴⁵ Marshall McLuhan, *Understanding Media*, p. 7 (1964: 2000 ed. Routledge Classics) (emphasis added).

⁴⁶ Luciano Floridi website, Research page, <http://www.philosophyofinformation.net/research/>.

⁴⁷ Luciano Floridi, “The Philosophy of Information: Ten Years Later”, *Meta*, Manuscript No. 1647, p. 1 (2010).

⁴⁸ Luciano Floridi, *Information: A Very Short Introduction* (2010: Oxford University Press).

⁴⁹ Maury D. Shenk, “Informationology: A New Framework for Understanding the Roles of Digital Information”, *Privacy & Data Security Law Journal* (Nov. 2009), available at SSRN: <https://ssrn.com/abstract=1448542>.

In summary, the reason for our brief exploration of these points of view on information is to inform our analysis of permission frameworks for information markets. Distilling the views of Shannon, Innis, McLuhan, Floridi and related scholars to their essence, we believe that an effective permission framework for information markets must have three core attributes:

- **Precision** – ability to accurately convey permissions (*i.e.*, the core principle of Shannon’s information theory);
- **Breadth** – scope to convey any type of permission (combining the views of Shannon with the expansive point of view of McLuhan); and
- **Applicability** – comprehensibility and practicality of application in real-world markets and related interactions (drawing on the work of Innis, McLuhan and Floridi).

B. Permission Frameworks

To apply this three-part standard to designing a future permission framework, we begin with a consideration of existing permission frameworks.

Deontic Logic

An excellent starting place for analysis of permissions is the formal system of *deontic logic*, which is the branch of symbolic logic relating to permission, obligation and related concepts.⁵⁰ Deontic logic was first proposed by Austrian philosopher Ernst Mally in 1926,⁵¹ and substantially elaborated by Finnish philosopher Georg Henrik von Wright (who succeeded to the Cambridge philosophy professorship of Ludwig Wittgenstein).⁵²

⁵⁰ See “Deontic Logic”, *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/logic-deontic/>; “Deontic logic”, *Wikipedia*, https://en.wikipedia.org/wiki/Deontic_logic

⁵¹ “Mally’s Deontic Logic” *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/mally-deontic/>

⁵² See “Deontic Logic”, *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/logic-deontic/>; “Profession G H von Wright (obituary), *The Telegraph* (23 June 2003), <https://www.telegraph.co.uk/news/obituaries/1433783/Professor-G-H-von-Wright.html>

On its face, deontic logic satisfies our *breadth* requirement, because it covers permissions generally and extends further to obligations and related concepts such as prohibitions. Indeed, while we generally focus on permissions in this report, it would be appropriate to adopt a permission framework that covers related concepts like obligation and prohibition. This is the case both because concepts like obligation and permission are relevant to markets, and because these various concepts are not easily separable. For example, the permission “you may park here for an hour” both implies a quasi-contractual obligation not to park for more than an hour and a description of a regulatory obligation (or prohibition) not to park in a designated spot without appropriate permission.⁵³ Deontic logic also appears to satisfy our *precision* requirement, because it is expressed with mathematical precision using standard symbols of propositional logic. For example, the two key axioms that deontic logic adds to standard propositional logic are:

$$\begin{aligned} O(A \rightarrow B) &\rightarrow (OA \rightarrow OB) \\ PA &\rightarrow \neg O\neg A \end{aligned}$$

In English, these are respectively “if it ought to be that A implies B, then if it ought to be that A, it ought to be that B” and “If A is permissible, then it is not the case that it ought not to be that A”.⁵⁴ Or to take the more practical example above of car parking, if proposition A is “you park here”, then ‘PA’ means “you may park here” and ‘O¬A’ means “you ought not park here” (or “you are forbidden to park here”). This type of precision is important for making a permissions framework computable, such as via a Smart Ledger. (Further examples of the use of deontic logic to express permissions are in Appendix 2 of this report. And for those who are interested in the details of deontic logic, the online *Stanford Encyclopedia of Philosophy* provides a detailed but clear explanation of the essentials of deontic logic.⁵⁵)

⁵³ “Deontic Logic”, *Stanford Encyclopedia of Philosophy*, section 4.1, <https://plato.stanford.edu/entries/logic-deontic/#4.1>

⁵⁴ “Deontic logic: Standard deontic logic”, *Wikipedia*, https://en.wikipedia.org/wiki/Deontic_logic#Standard_deontic_logic

⁵⁵ See “Deontic Logic”, *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/logic-deontic/>

Such precision may not imply consistency with existing computing systems, however, since computers are generally based on predicate logic, which involves ‘predicates’ (*i.e.*, propositions that are generally either true or false). Traditional teaching of logic suggests that you cannot mix predicate “is” logic with deontic “ought” logic, as one of the authors of this report has pointed out in a preparatory paper.⁵⁶

There is no obvious reason that a computer cannot handle permissions and obligations, and indeed many existing computer systems do so to some extent: returning to our parking example, consider online systems that issue parking permits, or process parking tickets. However, these systems perform mechanical tasks that do not reflect a human-like understanding of permission and obligation, and rendering these concepts computable may be increasingly difficult in more complex situations. The same preparatory report mentioned in the previous paragraph suggests that “information permission and obligation questions might, in essence, be a combination of consequential issues (weighing up possible benefits against possible harm) and normative values (judgements pertaining to autonomy and justice)”.⁵⁷ Unfortunately, computers tend to be ineffective at both of these tasks (particularly the latter).

Furthermore, the precision of deontic logic also limits its applicability, since it is not easy (other than for an experienced student of logic) to link such mathematical propositions to practical situations. The *Stanford Encyclopedia of Philosophy* expresses the need to make this connection, without proposing a clear way to do so:

“[D]espite the fact that we need to be cautious about making too easy a link between deontic logic and practicality, many of the notions listed *are* typically employed in attempting to regulate and coordinate our lives together (but also to evaluate states of affairs). For these reasons, deontic logics often directly involve topics of considerable practical significance

⁵⁶ Ian Harris & Professor Michael Mainelli, “Permissions In A Distributed World: Outline Research Into Pragmatic Models”, Long Finance Pamphleteers paper (November 2017), <https://www.longfinance.net/news/blogs/pamphleteers/permissions-distributed-information-world-outline-research-pragmatic-models/>

⁵⁷ Harris & Mainelli (note above).

such as morality, law, social and business organizations (their norms, as well as their normative constitution), and security systems. To that extent, studying the logic of notions with such *practical* significance adds some practical significance to deontic logic itself.”⁵⁸

So deontic logic fails to satisfy our third requirement of *applicability*. In sum, using deontic logic as the basis for a useful permissions framework would require (a) ensuring that deontic propositions are generally computable and (b) robustly linking such propositions to practical instances of permission. Both of these are significant areas for further research.

Such research is also occurring in related computing fields. A related space to deontic logic is the Vienna Development Method:

“Vienna Development Method (VDM) is one of the longest-established formal methods for the development of computer-based systems. Originating in work done at the IBM Laboratory Vienna in the 1970s, it has grown to include a group of techniques and tools based on a formal specification language — the VDM Specification Language (VDM-SL). It has an extended form, VDM++, which supports the modeling of object-oriented and concurrent systems. Support for VDM includes commercial and academic tools for analyzing models, including support for testing and proving properties of models and generating program code from validated VDM models. There is a history of industrial usage of VDM and its tools and a growing body of research in the formalism has led to notable contributions to the engineering of critical systems, compilers, concurrent systems and in logic for computer science.”⁵⁹

VDM approaches provide a structure and a set of mappings with operators. Deontic logic might be usefully ‘compiled’ into working systems using such approaches.

⁵⁸ “Deontic Logic”, *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/logic-deontic/>

⁵⁹ “Vienna Development Method”, Wikipedia (downloaded 5 November 2018) - https://en.wikipedia.org/wiki/Vienna_Development_Method

Access Control

Since the primary failing of deontic logic under our three-part standard is its lack of applicability, it makes sense to turn to more practical permission frameworks, especially those already in use in information markets.

If one asks a computer developer what she uses as a permission framework, she is likely to offer the concept of access control for computer systems:

“Access control is a security technique that can be used to regulate who or what can view or use resources in a computing environment.

There are two main types of access control: physical and logical. Physical access control limits access to campuses, buildings, rooms and physical IT assets. Logical access [control] limits connections to computer networks, system files and data.”⁶⁰

Here, we focus on logical access control, although physical access control is also relevant to markets. All of us are familiar with logical access control in terms of the need to have a password when accessing a computing device or website, as well as sometimes providing an additional security credential such as a regularly changing access code or a biometric credential (*e.g.*, a fingerprint or facial image). Logical access control can also take many more granular forms, including:

- **access control list (ACL)** – whoever (or whatever – *e.g.*, a device on the ‘Internet of things’) is on the list can access a system or a particular file, object or other resource within the system;
- **role-based access control (RBAC)** – expands the granularity of control by restricting what operations can be performed on a resource, depending upon the role(s) of the user accessing it;
- **attribute-based access control (ABAC)** – further expands the granularity of control to consider factors such as user attributes, resource attributes and context (*e.g.*, time, location, IP address).⁶¹

⁶⁰ “Access control”, *TechTarget SearchSecurity*,

<https://searchsecurity.techtarget.com/definition/access-control>

⁶¹ “Role-based access control”, *Wikipedia*, https://en.wikipedia.org/wiki/Role-based_access_control

Another approach to access control is to provide ‘registries’ that track everything. The Coalition Of Automated Legal Applications (COALA) has produced some explorations of how Smart Ledgers fit into digital rights by placing registries on the Smart Ledger or blockchain:

“Blockchain technology moves us toward a solution to both of these problems [payments and scale]. It allows:

1. Secure content registries tying creators and works;
2. Reliable decentralized content repositories that cannot lose information and are not vulnerable to censorship by authorities of any kind;
3. Micropayments to creators for every use of their works;
4. Automated smart contracts for sales, licensing, and novel uses of works; and
5. Entirely new forms of collaboration and creation that allow people who do not know or trust each other to work together.”⁶²

There are various other flavours of logical access control. But notwithstanding this alphabet soup of access control methods, all of these methods are generally limited to access to information systems and resources on them. They have little or nothing to say about permissions on financial markets, health, or transportation, or anywhere else where the data in the domain has meaning (except of course regarding access to information systems used in these contexts).

That is, access control permission frameworks manifestly fail our *breadth* requirement. Nevertheless, access control systems do well in satisfying the requirement of *precision*, and often do well on *applicability* (at least within their scope), so they should not be dismissed in providing a model for a broader permission framework.

⁶² COALA, “How Blockchains Can Support, Complement, Or Supplement Intellectual Property”, Working Draft (May 2016) - <file:///C:/Users/Michael/Downloads/COALA%20IP%20Report%20-%20May%202016.pdf>

Differential Privacy

Another, more-recent permission framework, just recently coming into widespread use, is ‘differential privacy’, which allows access to information in a database without disclosing whether or how that information relates to an identifiable individual. For example, the technique can provide data on preferences for a consumer product across a group, without disclosing individual ‘likes’. The technique was first disclosed in a 2005 United States patent application by Microsoft and its inventors Cynthia Dwork (a Harvard computer scientist) and Frank McSherry (a Microsoft researcher).⁶³ It has some application similarities to zero-knowledge proofs, by which one party can prove to another party that they know a value x , without conveying any information apart from the fact that they know the value x .

Differential privacy received a significant boost in market acceptance in 2016 when Apple announced that it would deploy the technology as part of the iOS 10 update of its iPhone/iPad operating system, protecting (even from Apple) access to private individual data of Apple users.⁶⁴ However, there has been some controversy over researchers’ suggestions that Apple has not using the technology in a suitably privacy-protective manner, by choosing to insert less noise into data than might be necessary to mask individual identities.⁶⁵ Specifically, researchers demonstrated that Apple had set the differential

⁶³ “Differential data privacy”, U.S. patent no. US7698250B2 (13 April 2010), <https://patents.google.com/patent/US7698250B2/>; “Cynthia Dwork”, *Wikipedia*, https://en.wikipedia.org/wiki/Cynthia_Dwork; “Frank McSherry: Some Background”, <http://www.frankmcsherry.org/about/>

⁶⁴ See Andy Greenberg, “Apple’s ‘Differential Privacy’ Is About Collecting Your Data, But Not Collecting Your Data”, *Wired* (13 June 2016), <https://www.wired.com/2016/06/apples-differential-privacy-collecting-data/>; Tom Simonite, “Apple’s New Privacy Technology May Pressure Competitors To Better Protect Our Data”, *MIT Technology Review* (3 August 2016), <https://www.technologyreview.com/s/602046/apples-new-privacy-technology-may-pressure-competitors-to-better-protect-our-data/>

⁶⁵ Jun Tang, Aleksandra Korlova, Xiaolong Bai, Xueqiang Wang & , Xiaofeng Wang, “Privacy Loss in Apple’s Implementation of Differential Privacy on MacOS 10.12”, arXiv, <https://arxiv.org/abs/1709.02753>; see also Andy Greenberg, “How One of Apple’s Key Privacy Safeguards Falls Short”, *Wired* (15 September 2017), <https://www.wired.com/story/apple-differential-privacy-shortcomings/>

privacy loss parameter *epsilon* too high on MacOS (and even higher on iOS) given the aggregate private loss resulting from frequent use of Apple devices.⁶⁶

This controversy nicely illustrates the need for a mathematically precise permission framework. When an individual grants access to personal data to Apple (or another provider) using differential privacy, it would be extremely useful for her system to be able to read to value of *epsilon* and present it in intelligible terms (or indicate that *epsilon* is concealed, as in Apple's case).⁶⁷ Indeed, differential privacy inventor Cynthia Dwork and others have proposed a United States National Epsilon Registry to disclose exactly this type of information.⁶⁸

Technology publisher Tim O'Reilly has recently made a similar, more general point:

“Disclosure and consent as currently practiced are extraordinarily weak regulatory tools. They allow providers to cloak malicious intent in complex legal language that is rarely read, and if read, impossible to understand. Machine-readable disclosure similar to those designed by Creative Commons for expressing copyright intent would be a good step forward in building privacy-compliant services.”⁶⁹

As O'Reilly observes, a similar approach has worked in the domain of copyright. We believe that the same type of machine-readable and machine-computable precision would be useful for many other types of market permissions.

Differential privacy also performs reasonably well on *precision* and *applicability*, but lacks *breadth*. Unfortunately, this is a general phenomenon for modern permission frameworks. And there are good reasons for this. *Precision* and *applicability* are necessary for any permission to work well in our IT-enabled

⁶⁶ Tang et al. (see note above).

⁶⁷ See Tang et al. (see note above). The study was able to deduce *epsilon* values, even though Apple does not publicly disclose those values.

⁶⁸ Cynthia Dwork & George J. Pappas, “Privacy in Information-Rich Intelligent Infrastructure”, arXiv (June 6, 2017), <https://arxiv.org/abs/1706.01985>

⁶⁹ See Tim O'Reilly, *What's the Future and Why It's Up to Us* (2017: Random House Business Books), p. 180.

society. The market would simply not tolerate any permission framework that is regularly inaccurate, or impractical to implement.⁷⁰ However, *breadth* is a significant challenge, because it is not easy to generalise a permission framework across disparate domains without making it unwieldy, perhaps threatening its performance on the essential requirement of *applicability*.

These observations present serious issues for our three-part standard of *precision*, *breadth* and *applicability*, indicating that it could be unworkable, because of tensions among the three requirements. However, we do not believe that this is an insurmountable challenge, for reasons that follow.

C. A Way Forward

We could give many more examples of modern permission electronic permission frameworks (*e.g.*, systems for electronic ticketing, insurance claims, voting registration and actual voting), but we believe that the examples above are sufficient to illustrate the way towards a more general permission framework.

The most difficult problem to solve is the need to reconcile *precision* and *applicability* with *breadth*. Fortunately, there is an excellent model for doing so in the structure of modern computer operating systems and languages – *i.e.*, providing a core framework on which others can build domain-specific functions. There are many examples of the tremendous generative power of this approach, for example:

- the Linux operating system – Linus Torvalds spurred the future success of Linux in 1991 by releasing the ‘kernel’ of Linux that offers core system functions, based on those of the widespread, proprietary Unix operating system (but without proprietary code). Coupled with other open source components available from the Berkeley Software Distribution and elsewhere, this provided the basis for an operating system that has been

⁷⁰ A good example of this need for applicability is the lack of mass market uptake of the first widely-available strong encryption technique Pretty Good Privacy (PGP), largely due to usability issues. See Amit Katwala, “We’re calling it: PGP is dead”, *Wired UK* (May 17, 2018), <https://www.wired.co.uk/article/efail-pgp-vulnerability-outlook-thunderbird-smime>

elaborated extensively by others and now runs the majority of the world's computer servers.⁷¹

- Apple iOS and Android – A crucial component of the success of the iPhone, released by Apple in 2007, has been the ability of developers to build apps on Apple's iOS operating system and sell them on Apple's App Store. Google's competing Android operating system (itself built on Linux) and its Google Play store have taken the same approach.⁷² These products and the many thousands of apps available on them have changed our society so significantly that it has started to become rather difficult to remember how we functioned just over a decade ago, before they appeared.
- Python and AI programming – There are many programming languages, but Python is the one showing by far the most growth in recent years, and is approaching overall dominance, particularly in the developed world.⁷³ A major reason (probably *the* major reason) for this growth is that interest in artificial intelligence is exploding, and Python has become the most useful language for AI programming for a variety of reasons, notably that many of the best AI programming libraries (notably Google TensorFlow) are written for Python.⁷⁴

⁷¹ See Klint Finley, "Linux Took Over the Web. Now It's Taking Over the World", *Wired* (25 August 2016), <https://www.wired.com/2016/08/linux-took-web-now-taking-world/>; "History of Linux", *Wikipedia*, https://en.wikipedia.org/wiki/History_of_Linux

⁷² See generally Walter Isaacson, *Steve Jobs* (2011: Simon & Schuster).

⁷³ See David Robinson, "The Incredible Growth of Python", StackOverflow blog (6 September 2017), <https://stackoverflow.blog/2017/09/06/incredible-growth-python/>

⁷⁴ See "6 Reasons: Why Choose Python for AI Projects?", NewGenApps blog (29 May 2017), <https://www.newgenapps.com/blog/python-for-ai-artificial-intelligence-ml>; Jason Brownlee, "Introduction to the Python Deep Learning Library TensorFlow", *Machine Learning Mastery* (5 May 2016), <https://machinelearningmastery.com/introduction-python-deep-learning-library-tensorflow/>

We suggest that a similar approach can work for a permissions framework, illustrated in the following diagram, where the ‘user interface’ is at the top and the deeper levels at the bottom are fully-automated.

Logical Access Control	Physical Access Control	Privacy	Consumer Financial	Securities Trading	Travel	Government Services	E-commerce
Domain-Specific Permission Libraries							
Deontic Logic API							
Deontic Logic Translation Engine							
Underlying Computing Operating System (e.g., Linux, iOS, MacOS, Windows)							

Figure 2 - Structure of possible future permissions framework

This framework supports discrete permission frameworks for different market and societal domains, allowing for permission approaches to satisfy the *precision* and *applicability* requirements within each domain. Although permission approaches would likely to differ significantly between domains, the permission approach for each and every domain would rest upon a common application programming interface (API) based upon deontic logic. This API would be analogous to a programming language working across computing platforms. As discussed above, this API would need to solve the non-trivial problem of combining deontic logic of permission and obligation with traditional computer logic based upon propositions that are true or false. Finally, each computer operating system would need have a translation engine (analogous to an interpreter or compiler for a programming language), to allow the API to function on that operating system.

This is of course a very ambitious vision for permissions, and we do not expect it to emerge simply because we have suggested it in this report. What we do suggest is that such a framework could be a fruitful area for further research, with various important problems to be solved before it could be implemented.

3. Smart Ledgers: An Emerging Tool For Distributed Permissions

This brings us to the role of Smart Ledgers. A major part of the rationale for the Distributed Futures project, of which this report is a part, is that Smart Ledgers have a role to play in identity, documentation, and agreement exchange.^{75,76} A large number of Smart Ledger solutions are competing for market attention, ranging from decentralised open solutions like Ethereum⁷⁷ and Cardano⁷⁸ to private ledger solutions built by companies like IBM⁷⁹ and Microsoft.⁸⁰

Of course, there are other technical solutions that can handle permissioning. Centralised databases managed by a trusted entity have been and remain the most common approach for managing permissions. But Smart Ledgers can have major advantages for *distributed* permissions, especially that they are inherently distributed and open. Furthermore, Smart Ledgers are co-evolving with the information markets providing the opportunity for Smart Ledgers to offer permission solutions that are especially fit-for-purpose for such new markets.

For example, Tim O'Reilly, after making the observation quoted above about machine-readable privacy disclosure, emphasised the potential role of Smart Ledger solutions in implementing such disclosures:

“During the Obama administration, there was a concerted effort toward what is called ‘Smart Disclosure,’ defined as ‘the timely release of complex information and data in standardized, machine readable formats in ways that enable consumers to make informed decisions.’ New technology like

⁷⁵ Michael Mainelli, “Blockchain Could Help Us Reclaim Control of Our Personal Data”, Harvard Business Review, Harvard Business School Publishing Corporation (5 October 2017) -

<https://hbr.org/2017/10/smart-ledgers-can-help-us-reclaim-control-of-our-personal-data>

⁷⁶ Michael Mainelli, “Blockchain Will Help Us Prove Our Identities In A Digital World”, Harvard Business Review, Harvard Business School Publishing Corporation (16 March 2017) -

<https://hbr.org/2017/03/blockchain-will-help-us-prove-our-identities-in-a-digital-world>

⁷⁷ Ethereum Project home page, <https://www.ethereum.org/>

⁷⁸ IOHK | Cardano, <https://iohk.io/projects/cardano/>. The Cardano Foundation, one of the entities responsible for the Cardano project, is a core sponsor of the Distributed Futures project.

⁷⁹ IBM Blockchain, <https://www.ibm.com/blockchain>

⁸⁰ Microsoft Azure, Blockchain, <https://azure.microsoft.com/en-gb/solutions/blockchain/>

the blockchain can also encode contracts and rules, creating new kinds of ‘smart contracts.’ *A smart contracts approach to data privacy could be very powerful.* Rather than using brute force ‘Do Not Track’ tools in their browser, users could provide nuanced limits to the use of their data. Unlike paper disclosures, digital privacy contracts could be enforceable and trackable.”⁸¹

However, for Smart Ledgers to realise this potential as a tool for managing permissions, significant further work is needed, especially in two areas:

- they must have the **technical features** to manage the requirements of an evolved permission infrastructure; and
- they must satisfy the **legal requirements** of the countries in which they are used.

A. Technical Features

Technical requirements for Smart Ledgers as a permission infrastructure are both straightforward and complex. They are straightforward in the sense that they are primarily an engineering problem. There is no apparent reason that Smart Ledgers cannot provide the required technical functionality to compete with more-established centralised solutions. So-called “third generation” blockchain platforms are making substantial progress in providing such enhanced functionality, including enhanced scalability, interoperability, sustainability, privacy and governance.⁸²

But there is obviously also substantial complexity in designing a permissions infrastructure. As we noted in proposing a layered permissions framework in section 2 above, there are major areas for further research required to make such an approach workable. There will also be a variety of general IT problems to be solved for any future permissions infrastructure, such as security. The

⁸¹ See Tim O’Reilly, *What’s the Future and Why It’s Up to Us* (2017: Random House Business Books), p. 180.

⁸² See, e.g., Michael K. Spencer, “Third Generation Blockchains”, *Medium* (30 May 2018), https://medium.com/@Michael_Spencer/third-generation-blockchains-7d6137e3f78b; Sudhir Katwani, “Top Five Blockchain 3.0 To Watch Out For In 2018”, *CoinSutra* (18 April 2018), <https://coinsutra.com/3rd-generation-blockchain/>

approach of an immutable ledger provides significant security guarantees, but this does not mean that Smart Ledgers are without security issues.⁸³ For example, another Distributed Futures report examined the need to address future security issues associated with quantum computing and encryption.⁸⁴

Providing a full map of technical issues and required features for distributed permissions is well beyond the scope of this initial report, however, it is an area of research we believe to be fruitful.

B. Legal Requirements

Unlike technical features, the legal requirements for Smart Ledgers are not straightforward, for the main reason that they are evolving rapidly. Notwithstanding the legal uncertainties, it is our firm view that legal restrictions and regulation will not ultimately be a significant barrier to the adoption of Smart Ledgers for permissions. In fact, some regulation such as GDPR might favour the use of Smart Ledgers for several use cases.

Why are we confident that regulation will not stifle Smart Ledgers? Because this story has played out many times before. New technologies are usually initially unregulated, and then face increasing regulation as they begin to disrupt markets. That regulation can have major implications as to who benefits from a new technology, and how much they benefit – but it rarely if ever materially deprives society of access to the technology. The most recent prominent example is the Internet.

The Internet was initially largely unregulated, but online regulation has been a steadily rising tide, now focusing, for example, on issues such as “net neutrality”.⁸⁵ From a user perspective, however, the availability of innovative

⁸³ See, e.g., “Q: How secure is a blockchain, really? A: It turns out ‘secure’ is a funny word to pin down.”, *MIT Technology Review*, May/June 2018, p. 40.

⁸⁴ Maury Shenk, “The Quantum Countdown: Quantum Computing And The Future Of Smart Ledger Encryption”, Long Finance / Distributed Futures (February 2018), <https://www.longfinance.net/publications/long-finance-reports/the-quantum-countdown-quantum-computing-and-the-future-of-smart-ledger-encryption/>

⁸⁵ See, e.g., Everett Ehrlich, “A Brief History of Internet Regulation”, Progressive Policy Institute report (13 March 2014), <http://www.progressivepolicy.org/issues/economy/a-brief-history-of-internet-regulation-2/>

services on the Internet continues to multiply. Very few services have been killed by regulation (although many have failed commercially), with the most notable exceptions being entities that traded without authorisation in the intellectual property rights of third parties, like Napster.

In the remainder of this section, we consider how evolving legal regulation is likely to affect Smart Ledgers in the European Union, the United States, China and India.

European Union

The European Union has led the headlines on possible regulatory barriers to adoption of Smart Ledgers, because of the entry into force of the EU General Data Protection Regulation (GDPR) in May 2018. There has been widespread speculation that GDPR may be inconsistent with blockchain and Smart Ledgers, most notably because of tension between the inherent immutability of Smart Ledgers and GDPR requirements including the 'right to be forgotten'.⁸⁶

We disagree that this is a serious problem, although it does require serious thought. In short, there are clear ways in which GDPR can be interpreted in a manner that allows Smart Ledgers applications, and various technical approaches that maximise compliance of Smart Ledger applications with GDPR. We provide a detailed discussion of the main legal issues for Smart Ledgers and GDPR in Appendix 1 of this report, and a summary of those issues is in the box below.

⁸⁶ See, e.g., Blockchain Bundesverband (the German Blockchain Association), "Blockchain, data protection, and the GDPR", p. 2 (25 May 2018), ("GDPR was created before Blockchain and is already outdated, since it doesn't account for decentralized technologies."), https://www.bundesblock.de/wp-content/uploads/2018/05/GDPR_Position_Paper_v1.0.pdf; Anne Toth, "Will GDPR block Blockchain?", World Economic Forum Industry Strategy Meeting (24 May 2018), <https://www.weforum.org/agenda/2018/05/will-gdpr-block-blockchain/>; Andries Van Humbeeck, "The Blockchain-GDPR Paradox", wearetheledger blog, *Medium* (21 November 2017), <https://medium.com/wearetheledger/the-blockchain-gdpr-paradox-fc51e663d047>

But Doesn't GDPR Make Smart Ledgers Illegal? – Q&A

- Q. Is permanent/immutable storage of data on Smart Ledgers prohibited by GDPR?
- A. GDPR does not absolutely bar permanent storage of data, and there are ways to implement effective 'erasure' of Smart Ledger data through use of encryption and off-ledger storage.
- Q. Does the repeated processing of Smart Ledger data for transaction processing violate GDPR?
- A. Repeated processing is permitted in various circumstances, and newer distributed ledger protocols substantially reduce repeated processing.
- Q. How can one identify data controllers and data processors in the Smart Ledger context?
- A. It is not always easy, but it is not in principle harder than in many other data protection contexts.
- Q. Do Smart Ledgers require automated processing that is restricted by GDPR?
- A. Possibly, in some cases. But like the previous issue, this is one that comes up in many data protection contexts. And there are important exceptions to the restrictions on automated processing.

At a high level, the feared (but avoidable) collision between GDPR and Smart Ledgers under GDPR is a result of the highly-regulatory approach of EU legislators to matters involving personal data and individual rights. EU authorities have explicitly stated that they intend GDPR to be a “gold standard” for global data protection regulation.⁸⁷ This approach has led to detailed requirements under GDPR for any data-driven business or application – not just Smart Ledgers and other blockchain applications – and businesses and entities are finding ways to comply. EU authorities are explicitly pro-technology,⁸⁸ and

⁸⁷ European Data Protection Supervisor, “The History of the General Data Protection Regulation”, https://edps.europa.eu/data-protection/data-protection/legislation/history-general-data-protection-regulation_en (“The EU's data protection laws have long been regarded as a gold standard all over the world.”).

⁸⁸ See European Commission, “Technology and innovation”, <https://ec.europa.eu/energy/en/topics/technology-and-innovation>

we expect this attitude (together with legal approaches along the lines set out in Appendix 1) to help avoid serious legal pitfalls for Smart Ledgers in Europe.

United States

The regulatory approach of the United States is very different from that of the EU. It is the country of Silicon Valley, Amazon, Facebook, and Google, and technology businesses have thrived there in significant part because of the *laissez faire* attitude of government towards industry and technological innovation. This has been all the more so during the current administration, although it remains to be seen whether the swing of United States politics towards the right will be durable.

In any event, in the United States the market and not regulation is the primary arbiter of the fate of new businesses and technologies. Although there is increasing interest in privacy regulation in the United States – for example, California enacted a significant new privacy law in June 2018⁸⁹ – it is highly unlikely that United States privacy regulation will present significant issues for Smart Ledgers that are beyond those of GDPR.

China

China is an extremely capitalist economy, rivalling the United States in terms of overall lack of regulation of economic initiative. Increasing liberalisation of the Chinese economy (begun by Deng Xiaoping in the late 1970s) has been the central reason for the huge economic success of China over the last four decades.⁹⁰ China is now even beginning to overtake the United States in core

⁸⁹ See Art Neill, “What You Should Know About The New California Consumer Privacy Law”, *Forbes* (29 June 2018), <https://www.forbes.com/sites/artneill/2018/06/29/what-you-should-know-about-the-new-california-consumer-privacy-law>

⁹⁰ See Ezra F. Vogel, *Deng Xiaoping and the Transformation of China* (2011: Harvard University Press).

aspects of technology development, such as venture capital investment⁹¹ and artificial intelligence research and innovation.⁹²

However, China remains a country where state control by the Communist Party is crucial, and targeted controls in the technology sector remain very significant. For example, the ‘Great Firewall of China’ – which blocks or limits access from China to many foreign Internet addresses – has been a key reason for the success of domestic Chinese companies like Alibaba, Tencent and Baidu over foreign players like Amazon, Google and Facebook.⁹³

In the blockchain sector, China has taken major regulatory action by banning cryptocurrency trading and initial coin offerings.⁹⁴ However, this has not been associated with significant steps to restrict other uses of blockchain and Smart Ledgers, and we do not expect China to impose material restrictions on Smart Ledgers in the future unless they present major threats to Chinese economic or social order (as is feared with cryptocurrency trading). Nevertheless, it is certainly possible that China will regulate Smart Ledgers in a way that favours domestic competitors. Ultimately, the Communist Party will decide.

China may also affect the evolution of permissions infrastructure through international standards processes, which are in essence a form of regulation. Over the past 10 to 15 years, China has been moving from an approach of

⁹¹ See Yingzhi Yang, “China surpasses North America in attracting venture capital funding for first time as investors chase 1.4 billion consumers”, *South China Morning Post* (5 July 2018), <https://www.scmp.com/tech/article/2153798/china-surpasses-north-america-attracting-venture-capital-funding-first-time>

⁹² See Louis Lucas & Richard Waters, “The AI arms race: China and United States compete to dominate big data”, *Financial Times* (1 May 2018), <https://www.ft.com/content/e33a6994-447e-11e8-93cf-67ac3a6482fd>

⁹³ See Emily Rauhala & Elizabeth Dwoskin, “U.S. companies want to play China’s game. They just can’t win it.”, *The Washington Post* (22 December 2016), https://www.washingtonpost.com/world/asia_pacific/us-companies-want-to-play-chinas-game-they-just-cant-win-it/2016/12/22/0fffa35a-b7f3-11e6-939c-91749443c5e5_story.html

⁹⁴ See Joseph Young, “Despite Crackdown on Trading, Crypto and Blockchain in China Are Alive”, *Cointelegraph* (7 March 2018), <https://cointelegraph.com/news/despite-crackdown-on-trading-crypto-and-blockchain-in-china-are-alive>

adopting its own technology standards,⁹⁵ to one of using its considerable market clout to influence global standards.⁹⁶ Given the importance of borderless, global permissions for information markets, Chinese involvement and influence will almost certainly be crucial to any future global permissions standard that is adopted.

India

India is also a country with extensive state control. However, this takes a very different form than in China. Indian state control is typically through extensive and detailed national and regional regulation (sometimes known as the ‘Licence Raj’⁹⁷), in contrast to the Chinese model of centrally-driven and often rather vague (or even unpublished) regulation.

The power of Indian bureaucracy can be very significant for technology companies. For example, the Telecom Regulatory Authority of India ruled in 2016 that Facebook was not permitted to offer its Free Basics service in India because of concerns regarding Internet neutrality.⁹⁸ Likewise, decisions of the Indian bureaucracy appear likely to have similarly significant effects on the future use of Smart Ledgers in the country.

A particularly important factor affecting whether Smart Ledgers will have a significant impact on permissions in India is the existing role of Aadhaar,⁹⁹ a centralised, state-run biometric ID system that as of early 2018 covered nearly 90% of the Indian population and is used for a wide variety of permissioning applications. Given this huge state investment in centrally-managed

⁹⁵ See, e.g., Richard Shim, “China implements new Wi-Fi security standard”, *CNet* (2 December 2003), <https://www.cnet.com/news/china-implements-new-wi-fi-security-standard/>; “Made-in-China WLAN standard”, Certified Wireless Network Professionals (8 May 2010), <https://www.cwnp.com/forums/posts?Made-in-China-WLAN-standard-123115>

⁹⁶ See Dave Burstein, “China: We Lead 3GPP Wireless Standards”, CircleID blog (26 May 2018), http://www.circleid.com/posts/20180526_china_we_lead_3gpp_wireless_standards/

⁹⁷ “Licence Raj”, *Wikipedia*, https://en.wikipedia.org/wiki/Licence_Raj. The Licence Raj has been somewhat dismantled since 1990.

⁹⁸ Rahul Bhatia, “The inside story of Facebook’s biggest setback”, *The Guardian* (12 May 2016), <https://www.theguardian.com/technology/2016/may/12/facebook-free-basics-india-zuckerberg>

⁹⁹ Unique Identification Authority of India (Aadhaar), <https://www.uidai.gov.in/>

permissions, it may be difficult for Smart Ledgers to play a meaningful role in permissions in India, in the medium-term at least.

Conclusion

Smart Ledgers have some inherent advantages over centralised solutions. They start off decentralised, providing advantages such as resilience and availability. They provide tremendous flexibility, perhaps too much. They provide excellent audit trails. But they are complex and harder to understand. A formal 'logic' or 'algebra' for permissions would help to ensure that the basic permissions an individual seeks over their digital 'items' can be expressed and implemented and tested.

In turn, a 'logic' or 'algebra' will incur legal technicalities that will differ by jurisdiction, e.g. Europe, United States, China, or India. While these technicalities may lead to tensions, we have shown that the most rigorous case, GDPR, is compatible with Smart Ledger approaches.

The criteria for successful permissioning systems appear to be:

- **Precision** – ability to accurately convey permissions;
- **Breadth** – scope to convey any type of permission;
- **Applicability** – comprehensibility and practicality of application in real-world markets and related interactions.

We anticipate a blend of access control, registries, and other traditional computer techniques, combining with new technologies such as differential privacy and zero-knowledge proofs, to provide a rich toolset for permissions. Our suspicion is that the toolset may be too rich and complex to be sure that it's doing what it was intended to do. Thus, a genuine area for research is figuring out how a formal deontic logic of 'may' and 'ought' can be implemented at the heart of Smart Ledgers to provide intrinsic permission structures similar to the way in which Smart Ledgers provide intrinsic timestamping and audit trails. Such a deontic logic might also help simplify matters.

There will be ongoing competition around the globe between centralised and distributed permission infrastructures for information and data markets. For Smart Ledgers to play a significant role (which we believe is possible), it is important for the issues set out in this report to be further researched and addressed by Smart Ledger proponents and developers.

Appendix 1 - Distributed Ledgers & GDPR

At the same time as Smart Ledgers and distributed ledger technology generally have been gathering public attention, the EU has been transitioning to a new data protection framework under GDPR, which aims to update EU data protection laws to address new technologies. GDPR was adopted in 2016 and took effect on May 25, 2018, and because distributed ledger technology has been advancing so rapidly, the drafters of GDPR did not appear to take this technology into account.¹⁰⁰ The result is a potential disconnect between GDPR framework and Smart Ledgers.

Although such issues require careful analysis, we believe that there are clear ways through the challenges for Smart Ledgers. In this appendix, we address four principal GDPR challenges that have been identified for distributed ledgers:

- the **permanent availability** of data on distributed ledgers to all who have access to the ledger;
- the **widespread and ongoing processing of ledger data**;
- the potential difficulty of **identifying data controllers** (*i.e.*, persons or entities responsible for data processing under GDPR) **and data processors** (*i.e.*, persons or entities that process data on behalf of a data controller) in the context of a distributed ledger; and
- the potential use of distributed ledgers for **automated decision-making**.

The first two issues are particular challenges for distributed ledgers and relate directly to the use of an immutable, distributed ledger. The latter two issues apply in a variety of data protection contexts but present some specific issues for distributed ledgers. We consider each of the four issues in turn below.

1. Data Permanence

The most frequently identified tension between GDPR and distributed ledger technology involves the inherent immutability of distributed ledgers.

¹⁰⁰ See, e.g., Blockchain Bundesverband (the German Blockchain Association), “Blockchain, data protection, and the GDPR”, p. 2 (25 May 2018), (“GDPR was created before Blockchain and is already outdated, since it doesn't account for decentralized technologies.”), https://www.bundesblock.de/wp-content/uploads/2018/05/GDPR_Position_Paper_v1.0.pdf

Immutability is a key feature of distributed ledgers, enhancing their security and allowing them to provide a permanent, verifiable, public record of transactions.

However, this immutability presents a potential for conflict with GDPR principles, especially:

- the **storage limitation principle** under GDPR Art 5(1)(e) that “personal data shall be ... kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed”; and
- the **right to erasure** (also known as the ‘right to be forgotten’) under GDPR Art 17, which provides an obligation on data controllers to erase data under specified circumstances.

As noted above, the main reason for the tension between these principles and distributed ledger immutability appears to be, simply, that the authors of GDPR did not anticipate distributed ledgers at the time that GDPR was adopted. But this does not mean that GDPR prevents or seriously impedes the deployment of these technologies, for two main reasons.

First, neither of the above principles is absolute. The storage limitation principle only restricts ongoing storage for “longer than is necessary for the purposes for which the personal data are processed”. Since immutability is a fundamental feature of distributed ledgers, it follows that with adequate advance notice of these functions users of the technology can be considered to have accepted that use inherently involves permanent storage and that this is “necessary for the purposes for which the personal data are processed”.

Likewise, the right to erasure applies only in specified circumstances – most importantly where (a) the “personal data are no longer necessary in relation to the purposes for which they were collected or otherwise processed”, (b) the data subject has withdrawn consent (if processing is based upon consent) or (c) the data controller processes personal data based on its “legitimate interests” without adequate justification.¹⁰¹ Where permanent storage is *required*, as in

¹⁰¹ GDPR Art. 17(1)(a)-(c). The other circumstances giving rise to a right of erasure involve unlawful processing, legal requirements for deletion and data regarding children. GDPR Art. 17(1)(d)-(f).

the case of distributed ledgers, there are likely to be relatively few circumstances in which these bases for erasure would be applicable. Among other things, processing in the distributed ledger context is frequently justifiable as necessary to perform a contract with the user¹⁰² (in which case withdrawal of consent is likely to be irrelevant).

Second, it is possible to design a distributed ledger to allow for the effective erasure of personal data, by storing that data off the ledger itself, or alternatively on the ledger in encrypted form that can only be decrypted by the holder of a private key. These approaches allow the data in question to be ‘erased’ by (a) deleting pointers to off-ledger data and/or (b) deleting private keys used for storing data either on-ledger or off-ledger (this latter options requires private keys to be assigned in a sufficiently granular fashion – *e.g.*, for particular data items or sets of data).¹⁰³ Indeed, some data protection authorities have already concluded that irreversible encryption constitutes erasure,¹⁰⁴ and both of these methods can involve irreversible encryption (the latter always does, and the former does if off-ledger data is encrypted).

2. Widespread And Ongoing Processing Of Ledger

Although distributed ledgers have been used for decades,¹⁰⁵ their recent explosive growth was initially driven by the Bitcoin protocol, which allows trust-free transaction verification through a proof of work consensus protocol.¹⁰⁶ This protocol requires every node of the Bitcoin network that wishes to engage in ‘mining’ of new bitcoins to repeatedly process new blocks of transactions (each time with a different ‘nonce’, or padding data) using a hash algorithm, until the calculated hash is below a certain value. Some have argued that this repeated processing – to the extent the processed blocks include personal data – conflict with GDPR principles, especially:

¹⁰² GDPR Art. 6(1)(b).

¹⁰³ Andries Van Humbeeck, “The Blockchain-GDPR Paradox”, *wearetheledger* blog, *Medium* (21 November 2017), <https://medium.com/wearetheledger/the-blockchain-gdpr-paradox-fc51e663d047>

¹⁰⁴ See Hogan Lovells, “A guide to blockchain and data protection”, p. 15 (September 2017), https://www.hलगage.com/_uploads/downloads/5425GuidetoblockchainV9FORWEB.pdf

¹⁰⁵ See Arvind Narayanan & Jeremy Clark, “Bitcoin’s Academic Pedigree”, *ACM Queue* (29 August 2017), <https://queue.acm.org/detail.cfm?id=3136559> (“Bitcoin’s ledger data structure is borrowed, with minimal modifications, from a series of papers by Stuart Haber and Scott Stornetta written between 1990 and 1997 (their 1991 paper had another co-author, Dave Bayer).”).

¹⁰⁶ See Bitcoin Wiki, “Proof of work”, https://en.bitcoin.it/wiki/Proof_of_work

- the **data minimisation principle** under GDPR Art 5(1)(c) that “personal data shall be ... limited to what is necessary in relation to the purposes for which they are processed”; and
- the **right to restriction of processing** under GDPR Art. 18 and the **right to object to processing** under GDPR Art. 21, which require data controllers to terminate or restrict the processing of personal data upon request in certain circumstances.

The conflict between these provisions and distributed ledgers is less obvious than with respect to data permanence. Furthermore, similar to the case of data permanence, there are two main bases for addressing any legal concerns.

First, the principles themselves are subject to significant limitations. With respect to the data minimisation principle, there is a good argument that multiple hashing of personal data does not mean that such data are not “limited to what is necessary in relation to the purposes for which they are processed”. That is, multiple hashing does not increase the *amount of personal data* that is processed – which is the core focus of the data minimisation principle – but rather relates to the *number of times* that the data are processed.

Likewise, the right to restriction of processing and right to object to processing apply only in specified circumstances, which are narrower than those triggering the right to erasure, *i.e.*, where (a) there is a challenge to processing based upon “legitimate interests” (as for the right to erasure), (b) there is a challenge to accuracy of personal data, (c) processing is unlawful, (d) the data controller no longer needs the data but the data subject (*i.e.*, the individual to whom the data relate) wishes the data to be retained for reasons related to legal claims, or (e) the processing involves use of profiling for direct marketing.¹⁰⁷ On a distributed ledger, there may be no way to entirely stop processing of the ledger in any of these circumstances; however, it is entirely possible to design distributed ledger solutions so that any personal data is encrypted and cannot be processed in a manner that discloses the data in these circumstances.

Second, not all distributed ledger protocols are created equal. For example, many newer distributed ledger protocols include consensus protocols that undertake significantly less frequent processing or confirmation of ledger

¹⁰⁷ GDPR Arts. 18(1), 21(1) & 21(2).

transactions than does a proof-of-work consensus protocol. This is the case, for example, for ‘proof of stake’ consensus protocols like those proposed by Cardano¹⁰⁸ and EOS¹⁰⁹ protocols, and planned for Ethereum.¹¹⁰ There are numerous proposals for alternative architectures where perhaps the only common elements binding them together as ‘Smart Ledgers’ are ‘immutability’ and ‘embedded computer code that can be executed at a future date’.

3. Identifying Data Controllers & Data Processors

Unlike the previous two issues, the challenge of identifying data controllers and data processors is not specific to distributed ledgers. GDPR defines ‘controller’ as:

“the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data”.¹¹¹

‘Processor’ is defined as:

“a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller”.¹¹²

Applying these definitions to complex, multi-party applications and technical ecosystems (consider, for example, the interactions of buyers, sellers and payment providers on a platform like Amazon or eBay) is a frequent challenge for data protection practitioners. However, despite potential ambiguities, it is our experience that a practical, good faith approach to defining controller and processor roles is generally effective from a regulatory perspective.

For example, the following approach appears sensible:

- users that store data and build applications on a distributed ledger are data controllers with respect to any personal data that they process or store;

¹⁰⁸ Ourobours Proof of Stake Algorithm, <https://cardanodocs.com/cardano/proof-of-stake/>

¹⁰⁹ Brady Dale, “EOS Is Coming, If Anyone Can Figure Out How To Vote”, *Coindesk* (30 May 2018), <https://www.coindesk.com/eos-coming-anyone-can-figure-vote/>

¹¹⁰ Shiraz Jagati, “Ethereum Proof of Stake Protocol Under Review”, *CryptoSlate* (22 April 2018), <https://cryptoslate.com/ethereums-proof-of-stake-protocol-in-review/>

¹¹¹ GDPR Art. 4(7).

¹¹² GDPR Art. 4(8).

- distributed ledger nodes are data processors when they process transactions for others; and
- the operator of a distributed ledger (*i.e.*, one that is not purely decentralised) is a data controller with respect to personal information of individuals with which it interacts in order to operate the ledger.

This approach to defining roles of data controllers and data processors is consistent with approaches recently recommended by the German Blockchain Association¹¹³ and others.

4. Automated Decision-Making

Like the previous issue, the question of automated decision-making is not specific to distributed ledgers. Article 22 of GDPR restricts automated decision-making without human involvement “which produces legal effects concerning [an individual] or similarly significantly affects him or her.” This provision has generated substantial interest and concern in the technology community because a wide variety of emerging applications – particularly those involving artificial intelligence and machine learning – use automated decision-making.¹¹⁴ For distributed ledgers, the Article 22 restriction is only relevant to a distributed ledger application to the extent that the application uses automated decision-making. Whether any application in fact does so must be assessed on a ledger-specific and application-specific basis. Furthermore, there are important exceptions to the Article 22 restriction, including where:

- there is a ‘human in the loop’ in some non-trivial respect – *i.e.*, decision-making is not purely automated;
- an automated decision does not produce “legal effects” or similar effects;
- the automated decision or processing “is necessary for entering into, or performance of, a contract between the data subject and a data controller” (GDPR Art. 22(2)(a)); or
- the automated decision is made with the data subject’s explicit consent.

¹¹³ See Blockchain Bundesverband (the German Blockchain Association), “Blockchain, data protection, and the GDPR”, pp. 5-7 (25 May 2018), (“GDPR was created before Blockchain and is already outdated, since it doesn't account for decentralized technologies.”),

https://www.bundesblock.de/wp-content/uploads/2018/05/GDPR_Position_Paper_v1.0.pdf

¹¹⁴ See, *e.g.*, Pomin Wu, “GDPR and its impacts on machine learning applications”, *Medium* (7 November 2017), <https://medium.com/trustableai/gdpr-and-its-impacts-on-machine-learning-applications-d5b5b0c3a815>

Appendix 2 - Permission Framework Examples

Translation From Human Language To Deontic Logic Propositions

	What a Human Hears	High-Level Proposition	Propositional Variables ¹¹⁵	Deontic Proposition
Identity Domain	You are an authorised user of this computer system	Person X may access resource R	$AR_X = X$ accesses resource R	$P(AR_X)$
	If you are in the finance department, you may access the accounting system	If person X belongs to group G, she may access resource R	$AR_X = X$ accesses resource R $G = \text{group } G$	$If X \in G \rightarrow P(AR_X)$
	Would Ms Jones please go to the ticketing desk	If recipient of message is person X, she should take action A	$U = \text{recipient of message}$ $A_X = X$ takes action A	$If U = X \rightarrow O(A_U)$ $If U = X \rightarrow O(A_X)$
	Sorry, no admittance for under 18s	If person X is under age K, she may not access resource R	$K_X = \text{age of } X^{116}$ $AR_X = X$ accesses resource R	$If K_X < 18 \rightarrow \neg P(AR_X)$
	No ID, no entry	If person X cannot prove she is over age K, she may not access resource R	$K_X = \text{age of } X$ $ID_X = \text{identification documents in } X\text{'s possession}$ $AR_X = X$ accesses resource R	$If (K_X > 18) \neg \vdash ID_X \rightarrow \neg P(AR_X)$
Health Domain	You ought to see the doctor	Person X ought to visit a doctor	$MD_X = X$ visits a doctor	$O(MD_X)$
	Your temperature is 40.2°! Call the doctor.	If person X's body temperature is over 40°C, she should see a doctor	$T_X = \text{body temperature of } X$ $MD_X = X$ visits the doctor	$If T_X > 40 \rightarrow O(MD_X)$
	Now that you're 50, you should have your blood pressure checked regularly	If person X is over age K, she should take action A	$K_X = \text{age of person } X$ $A_X = X$ takes action A ¹¹⁷	$If K_X > 50 \rightarrow O(A_X)$

¹¹⁵ For all propositions, X represents a human individual.

¹¹⁶ In a practical permissions system, such a variable would likely be generalised to accommodate multiple characteristics of person X.

¹¹⁷ In this example, the relevant action is handled with a general variable (A_X), while in the previous example the relevant action is handled with a variable specific to the action (MD_X).

Information Rules: Smart Ledger Architectures & Distributed Permissions

	You can look at my health data when I'm in an emergency ward and in a critical situation	Person X may access resource R	$AR_X = X$ accesses resource R	$P(AR_X)$
	You may access free NHS services	Person X may access resource R	$AR_X = X$ accesses resource R	$P(AR_X)$
	You are not an EU resident, so NHS services are not available for free	If person X's country of residence is not an EU member state, she may not access resource R	$C_X =$ country of residence of X EU = the member states of the EU $AR_X = X$ accesses resource R	$If C_X \notin EU \rightarrow \neg P(AR_X)$
Gambling Domain	This website is limited to over 18s	If person X is under age K, she may not access resource R	$K_X =$ age of person X $AR_X = X$ accesses resource R	$If K_X < 18 \rightarrow \neg P(AR_X)$
	Minimum table stake is £500	If person X acquires tokens exceeding value V, she may access resource R	$VT_X =$ value of tokens acquired by X $AR_X = X$ accesses resource R	$If VT_X \geq 500 \rightarrow P(AR_X)$
	This casino is for select patrons	If person X has assets exceeding value V, she may access resource R	$VA_X =$ value of assets owned by X TA = asset threshold for particular casino (e.g., £5 million) $AR_X = X$ accesses resource R	$If VA_X \geq TA \rightarrow P(AR_X)$
	Never hit to 17 when the dealer shows a six	In blackjack, one ought not request a card when the dealer's visible card is C	$H_X = X$ requests a card in blackjack $C_D =$ visible card(s) of dealer	$If CB_D = 6 \rightarrow \neg O(H_X)$
	You should not gamble	The recipient of a message should not gamble	U = recipient of message $G_X = X$ gambles	$\neg O(G_U)$
Music	This website is limited to music memberships	If person X is not a member, she may not access resource R	$M_X =$ membership of person X $AR_X = X$ accesses resource R	$If M_X = valid \rightarrow \neg P(AR_X)$
	Listening to this music costs \$0.005	If person X acquires tokens exceeding value V, she may access resource R	$VT_X =$ value of tokens acquired by X $AR_X = X$ accesses resource R	$If VT_X \geq 0.005 \rightarrow P(AR_X)$

Information Rules: Smart Ledger Architectures & Distributed Permissions

	You can listen to this music 20 times for each purchase	If person X has purchased (P) the music they have Y 'listens'	P_x = music purchase owned by X L_m = listens so far AR_x = X accesses resource R	$If P_x = valid \ \& \ L_m < 21 \rightarrow P(AR_x)$
	You can share this music to a depth of three people	If person X has purchased (P) the music they have H 'handons'	P_x = music purchase owned by X H_m = handons so far AR_x = X accesses resource R	$If P_x = valid \ \& \ H_m < 3 \rightarrow P(AR_x)$

Principal Authors



Maury Shenk

Managing Director, Lily Innovation

Maury's experience focuses at the intersection of technology, law and business. He is founder and managing director of Lily Innovation, through which he handles a portfolio of activities including private equity and corporate finance, legal advisory, directorships, start-up investing and teaching/writing. Maury is a dual-qualified US/UK lawyer and former managing partner of the London office of global law firm Steptoe & Johnson, where he remains an advisor; general counsel of China-based private equity fund Spring Capital Asia; and director of testing and certification company PeopleCert and recycling compliance company Valpak. He has a deep practical understanding of technology, especially IT and telecommunications, artificial intelligence, information security and green technology. Maury is a graduate of Harvard College and Stanford Law School. He is a lover of languages – a native speaker of English (the American version), proficient in French and Russian, comfortable in Mandarin Chinese and Spanish, and a dilettante in German, Italian, and Norwegian. He is also an avid competitive and recreational sailor.



Professor Michael Mainelli FCCA FCSI FBCS, Executive Chairman, Z/Yen Group

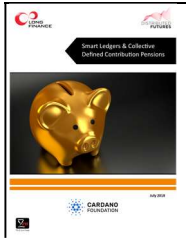


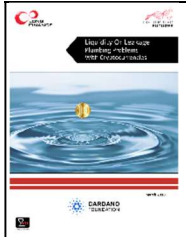


A qualified accountant, securities professional, computer specialist, and management consultant, educated at Harvard University and Trinity College Dublin, Michael gained his PhD at London School of Economics where he was also a Visiting Professor. He began his career as a research scientist, later becoming an accountancy-firm partner and a director of Ministry of Defence research. During a spell in merchant banking in 1994, he co-founded Z/Yen, the City of London's leading commercial think-tank. He has led Z/Yen from creating Smart Ledgers (aka blockchains) through the Financial Laboratory, Taskforce 2000, Long Finance, Global Financial Centres Index, Global Green Finance Index, and Global Intellectual Property Index. He is a non-executive director of two listed firms and a regulator, Emeritus Professor at Gresham College, Fellow of

Goodenough College, and a past Master of the Worshipful Company of World Traders. His third book, **The Price of Fish: A New Approach to Wicked Economics and Better Decisions**, co-authored with Ian Harris, won the 2012 Independent Publisher Book Awards Finance, Investment & Economics Gold Prize.

Acknowledgments

The authors would like to thank Ian Harris for his enormous contribution to the original thinking behind this report.

Other Selected Distributed Futures Publications

	Title	Authors	Year	Publisher
	Smart Ledgers & Collective Defined Contribution Pensions	Iain Clacher, Con Keating, and David McKee	2018	Long Finance (July 2018), 47 pages.
	Timestamping Smart Ledgers - Comparable, Universal, Traceable, Immune	Sam Carter	2018	Long Finance (June 2018), 55 pages.
	The Economic Impact Of Smart Ledgers On World Trade	Centre for Economics and Business Research	2018	The Worshipful Company of World Traders and Long Finance (April 2018) 78 pages.
	Liquidity Or Leakage - Plumbing Problems In Cryptocurrencies	Rodney Greene and Bob McDowall	2018	Long Finance (March 2018), 61 pages.
	Get Smart About Scandals: Past Lessons For Future Finance	Professor Tim Connell and Bob McDowall	2018	Long Finance (March 2018), 102 pages.
	The Quantum Countdown: Quantum Computing And The Future Of Smart Ledger Encryption	Maury Shenk	2018	Long Finance (February 2018), 61 pages.



Distributed Futures is a significant part of the Long Finance research programme managed by Z/Yen Group. The programme includes a wide variety of activities ranging from developing new technologies, proofs-of-concept demonstrators and pilots, through research papers and commissioned reports, events, seminars, lectures and online fora.

Distributed Futures topics include the social, technical, economic, and political implications of Smart Ledgers, such as identity, trade, artificial intelligence, cryptography, digital money, provenance, FinTech, RegTech, and the internet-of-things.

www.distributedfutures.net



Cardano Foundation is a Smart Ledger and cryptocurrency organisation based in Zug, Switzerland. The Foundation is dedicated to act as an objective, supervisory and educational body for the Cardano Protocol and its associated ecosystem and serve the Cardano community by creating an environment where advocates can aggregate and collaborate.

The Foundation aims to influence and progress the emerging commercial and legislative landscape for blockchain technology and cryptocurrencies. Its strategy is to pro-actively approach government and regulatory bodies and to form strategic partnerships with businesses, enterprises and other open-source projects. The Foundation's mission is the promotion of developments of new technologies and applications, especially in the field of new open and decentralised software architectures.

www.cardanofoundation.org



"When would we know our financial system is working?" is the question underlying Long Finance's goal to improve society's understanding and use of finance over the long term. Long Finance aims to:

- ◆ expand frontiers - developing methodologies to solve financial system problems;
- ◆ change systems - provide evidence-based examples of how financing methods work and don't work;
- ◆ deliver services - including conferences and training using collaborative tools;
- ◆ build communities - through meetings, networking and events.

www.longfinance.net



Z/Yen is the City of London's leading commercial think-tank, founded to promote societal advance through better finance and technology. Z/Yen 'asks, solves, and acts' on strategy, finance, systems, marketing and intelligence projects in a wide variety of fields. Z/Yen manages the Long Finance initiative.

Z/Yen Group Limited
41 Lothbury, London EC2R 7HG, United Kingdom
+44 (0) 207-562-9562 (telephone)
hub@zyen.com (email)
www.zyen.com