Dissertation

How I Learned to Stop Worrying and Love Blockchain

Implications and Applications of Blockchain

Pavan Katkar

This document was submitted as a dissertation in May August in partial fulfillment of the requirements of the doctoral degree in public policy analysis at the Pardee RAND Graduate School. The faculty committee that supervised and approved the dissertation consisted of C. Richard Neu (Chair), Angela O'Mahony, and Philip Evans.

Funding support was provided by the Lawrence Livermore National Laboratory.

RAND PARDEE RAND GRADUATE SCHOOL

For more information on this publication, visit http://www.rand.org/pubs/rgs_dissertations/RGSDA1084-1.html

Published 2021 by the RAND Corporation, Santa Monica, Calif.

RAND® is a registered trademark

Abstract

The technology underpinning contentious cryptocurrencies (like bitcoin) is widely known as blockchain. This dissertation explored the implications and applications of blockchain. By 'implications' I mean what kinds of policy problems is it enabling, and by 'applications' I mean what kinds of policy problems can we address using blockchain.

The first part of this dissertation explored some of the widely discussed policy problems associated with blockchain to find out that blockchain is not creating any new kinds of policy problems. However, it has the potential to scale up some specific law enforcement issues associated with transaction of illicit goods and services.

The second part of this dissertation conceptualizes a blockchain based ecosystem to demonstrate that this tool can be used in orchestrating an ecosystem geared towards mitigating large scale ($n \gg 1$) cyber attacks. This exploration led to the insight that if addressing a policy issue requires the alignment of economic incentives of multiple independent entities in a cooperative way, then blockchain can be quite an effective tool for that purpose.

Based on these analyses, this dissertation suggests that while blockchain (and cryptocurrencies) may be increasing the scale of some select law enforcement issues, there is no need for new policies as such; existing policies are robust enough to address those issues. Secondly, governmental and commercial organizations may benefit by exploring how to apply this tool in addressing complex policy problems that need coordination of multiple independent entities.

To my aunt...

Acknowledgments

I acknowledge that completing this dissertation was a truly eye-opening experience. It was an eye-opening experience because through this dissertation not only did I study blockchain, cryptocurrencies, and public policy but also, on a personal level, I explored – neither willingly nor unwillingly – some of the most fundamental questions of one's existence.

It was extremely surprising to me how this dissertation focused on studying the implications and applications of a new technology for public policy, mapped so well on to the "metaphysical" level. In that sense, I wrote two dissertations – one that is metaphysical, and the other that is not. The one that you are about to read is, of course, the technical version of my dissertation. That said, it is likely that if one were to explore this dissertation plainly while squinting their eyes in the "right places", then they will inevitably find the metaphysical version too. Of course, it goes without saying that it is completely coincidental that the technical version maps on to the metaphysical version so well. There was no deliberate effort on my part in bringing about that mapping. It is too difficult to explain how the technical maps on to the metaphysical level, so I will not even try.

When I began writing this dissertation, I had no idea how to intellectually penetrate a topic that I knew nothing about. I relied heavily on the wisdom of my dissertation committee to learn how to think through an unknown topic that was completely obscured by hype and speculation. Without them I would not have completed this dissertation. Hence, I would like to begin by thanking them first.

I thank the Chairperson of the committee, C Richard Neu (Sr. Economist at RAND), for guiding me in my quest to think deeply (and clearly) about something. He was in many ways the Ideal Chairperson. He ensured that I had enough freedom to explore whatever it is that I wanted to explore through this dissertation without letting anyone (including himself) impose their ideas on me. I could not have found a better Chairperson for my committee. I am deeply indebted to him.

Next, I would like to thank, Philip Evans (Senior Advisor at BCG's Henderson Institute), for showing me how to think broadly about a given topic. When the novice thinker in me learned how to think deeply, the excitement made it rather easy for me to lose track of the breadth of my thinking. Philip ensured that the breadth of my thinking did not suffer due to such excitement. Many of his comments were breadth-provoking and it would have been difficult for me to write a reasonably well-rounded dissertation without his guidance. I cannot thank him enough.

Next, I would like to thank, Angela O'Mahony (Assistant Dean for Academic Affairs at PRGS and Senior Political Scientist at RAND), for helping me in organizing and ordering my thinking. Every discussion I had with her resulted in greater clarity of presentation of ideas. Whenever I was stuck with and felt overwhelmed by the barrage of ideas I had, she swooped in and helped me cut through the clutter of my thinking and helped me organize them in a way that made it presentable and digestible (even for me). I have immense gratitude for her.

Apart from my committee members, I would like to thank Lawrence Livermore National Laboratory (LLNL), Project Air Force, and National Defense Research Institute for each providing a dissertation award that supported the completion of this dissertation. I would also like to specially thank Peter D Barnes (from LLNL), Ed Balkovich, and Richard Girven (both from RAND) for reading through the technical chapters and providing clarity-provoking comments on them. I would also like to thank Raffaele Vardavas (from RAND) for helping me understand complex systems in a mathematically rigorous way by hiring me to work along side him on projects focused on understanding human behavior.

Next, I would like to thank my girlfriend for teaching me one of the most important of skills: observation. She taught me to observe things in a way that separated the signal from the noise. Not only did I learn to understand and appreciate technically difficult topics but also philosophically oriented topics like art, and most importantly, love.

Next, I would like to thank my parents and my brother for explaining the basics of banking and money. I would also like to thank my colleagues at RAND and my friends (special shoutout to my buddy, Sachin Kodihalli) for introducing me to certain ideas at the right time that helped me get unstuck in my dissertation. In addition, I would like to thank an anonymous source for the support they provided.

Lastly, most of all I would like to thank my late aunt for teaching me to understand the mystery of life when she was alive that only made sense to me after she was gone. Obviously, any mistakes or discrepancies found in this journey are solely my responsibility and do not reflect anything about those who are acknowledged here.

– Pavan Katkar

Contents

List of Figures

1. Introduction

> "*All you need is the plan, the road map, and the courage to press on to your destination*"
>
> —Earl Nightingale

Blockchain is an innovative technology underpinning a form of currency called cryptocurrency – a digital currency with cryptographic properties that prevent counterfeiting. Ever since the development of blockchain along with the first cryptocurrency called bitcoin in 2009, there has been an exponential growth in the number of versions of blockchain, cryptocurrencies, and other related digital (crypto)assets. In some ways, an alternative financial ecosystem appears to be emerging that is centered around crypto financial assets.

A technological development that is not only enabling the creation of new currencies but also bootstrapping an alternative financial ecosystem can potentially create significant implications for public policy. Furthermore, blockchain has the potential to address some business issues like supply chain management and cross-border money transfer, and also to address some policy problems like healthcare information sharing and managing property rights.

This dissertation, therefore, is divided into two parts. The first part explores the potential policy implications of blockchain. It has four chapters. The first one is focused on introducing the concept of blockchain to a non-technical audience. Second one is focused on exploring the impact of blockchain on the role for intermediaries in various kinds of transactions. The third one focuses on exploring how blockchain creates value. Lastly, the fourth one focuses on discussing some of the major policy considerations that were identified in its preceding chapters.

The second part explores the potential for policy-relevant applications of blockchain. This section has two chapters. The first one focuses on identifying and discussing a policy relevant problem of large scale cyber attacks. The second chapter focuses on conceptualizing a blockchain based ecosystem to mitigate that issue.

At the end of this dissertation, an independent chapter is included to provide some of the main concluding thoughts on the whole phenomenon of blockchain. The remainder of this introductory chapter lays out more details of how this dissertation goes about exploring the implications and applications of this technology called blockchain.

1.1 Part 1: Implications of Blockchain

One of the important policy questions related to blockchain is: what should policy-makers do about this technology? Should they support it, oppose it, or ignore it? To address this overarching policy question there are multiple sub-questions that need to be investigated. First and foremost of them is: What is blockchain?

1.1.1 Chapter 2: What is Blockchain?

The chapter acts as a short primer on blockchain. It explores blockchain in detail to assist the reader in developing a generic understanding of what it does, how does it work, why are there so many implementations, and how is it being used. As a corollary, a working definition of blockchain is developed that sets the scope for the follow on discussion in the remaining chapters. The main aim of this chapter is to leave the reader with a reasonable understanding of what one may call the "blockchain phenomenon" without delving too much into the microscopic technical details of the technology. This chapter defines blockchain as *a technology that enables its users to create and exchange cryptographically protected value with each other.*

1.1.2 Chapter 3: Intermediaries

Since blockchain can enable its users to directly transact value with each other without the need for intermediaries, many blockchain enthusiasts claim that this may herald an age of disintermediation and decentralization, where financial institutions may not be necessary for people to transact with each other. This chapter explores the impact blockchain appears to be having on the role of intermediaries. One of the main insights that emerged through this exploration is that blockchain has the potential to disintermediate the *initiation-side* of the transaction but not necessarily the *completion-side* of the transaction. The chapter also explores how blockchain based ways of transacting are bringing in new kinds of intermediaries and what are some of the regulatory gaps that have been created due to these developments.

1.1.3 Chapter 4: Blockchain-Created Value

Since blockchain not just enables its users to exchange of value but also to create value, it felt important to explore the topic of blockchain-created value. This chapter closely examines value created through blockchain. It classifies blockchain-created value into two categories: a) value created by the use of blockchain as an instrument of record-keeping and market-making, and b) value created on blockchain in the form of cryptocurrencies. Both these categories are explored in detail to identify the main issues posed by blockchain-created value.

1.1.4 Chapter 5: Selected Policy Considerations

In this chapter, all the issues that were identified and flagged as important in the preceding chapters are explored in more detail from the perspective of what policy responses may be needed to address them. This chapter shows that most of the identified issues are related to blockchain being used for a) transacting illicit items, b) transacting in illicit ways, and c) transacting for illicit purposes. Examination of each of those categories led to the insight that although blockchain has the potential to aggravate these issues by increasing their scale, the existing policies are quite robust to address them. More than policy changes, law enforcement agencies may benefit from enhanced technological capabilities to scale up their investigative processes. The chapter also finds that policy makers in the US are, in fact, doing many important things to address these issues in an innovation-friendly manner.

1.2 Part 2: Applications of Blockchain

Since blockchain makes it easy to create a system of incentives for multiple entities to achieve predefined outcomes, this dissertation would be remiss to not conceptualize an application in reasonable detail to demonstrate the potential utility of blockchain. This section deliberately chooses a policy problem to demonstrate the potential role blockchain can play in addressing the problem.

1.2.1 Chapter 6: A Potential Application of Blockchain

This chapter first shows in detail that the main reason cyber attacks or Advanced Persistent Threats (APT) like WannaCry, Stuxnet, and others spread like wildfire is: identical software. Identical software implies that successfully hacking one system makes it significantly easy to hack other systems i.e. a successful zero-day[1] exploitation can rapidly evolve into a large-scale attack. In addition, the marginal cost of exploiting systems at a large scale quickly diminishes to zero, while the damages caused tend only to increase. This economic asymmetry makes these threats

[1] A vulnerability that is not yet known to the developers of the software and hence, there is no known solution to fix the vulnerability

"persistent" since hackers can break a system once and repeatedly use that method of exploitation to compromise other system. It can be significantly damaging, if not catastrophic, if adversaries compromised cyber-physical systems like unmanned aerial vehicles or autonomous vehicles in large numbers to cause massive accidents.

Based on recent research in computer science, this chapter suggests that one solution to this problem can be non-identical software with identical functionality. However, it also recognizes that implementing such a solution can be substantially complicated because of the challenges associated with logistics and supply chain of deploying software updates and third-party software applications (if needed). The chapter ends with making a case for why blockchain can be potentially useful in addressing these challenges.

1.2.2 Chapter 7: A Blockchained Ecosystem

This chapter takes a stab at conceptualizing a blockchain based ecosystem that can align the incentives of multiple entities in the software industry in such a manner that developing and deploying non-identical software with identical functionality can become sustainable. It borrows ideas from various blockchain-apps and distributed systems that are operational and applies them within the context of deploying non-identical software on different systems. It also conceptualizes a crypto-token called Eddy that can be used as a mechanism to align the incentives of the various entities in the ecosystem so that they behave in a cooperative manner.

The main insight that emerged from this exercise was that blockchain can be used as an effective tool when the problem being addressed has multiple independent entities with competing interests that may be blocking the resolution of the problem. Using blockchain and crypto-tokens it may be possible to incentivize the entities in the ecosystem to cooperate with each other in a way that can help in addressing the problem.

1.3 Chapter 8: Concluding Thoughts

The last chapter of this dissertation summarizes all the various trains of thought I followed throughout the process of writing this dissertation and tie them all together in a, hopefully, coherent manner. It is a forward looking chapter that delves deeper in the phenomenon of blockchain to identify what are some of the most fundamental changes blockchain can usher into our society if the technology is widely adopted.

2. Introduction to Blockchain

> "*When we look at something that is alien to us, that is beyond our comprehension, what do we see but ourselves?*"
>
> —Trevor Paglen

2.1 Origins of Blockchain

After the precipitation of the great financial crisis in 2008, there was a general sense of anger in the public toward financial institutions. [1] Many held the financial institutions responsible for the crisis because these institutions, the argument goes, managed financial risk inadequately. Towards the end of 2008, a document of mysterious origin was published online conceptualizing bitcoin – a peer-to-peer cash system (Nakamoto, 2008).

The paper stated that the bitcoin cash system would "allow online payments to be sent directly from one party to another without going through a financial institution." Furthermore, it stated that the current system where financial institutions act as trusted intermediaries that process electronic payments suffers from some inherent weaknesses that are associated with trust based models. These weaknesses, the paper claims, include: a) higher transaction costs due to the need for mediation by the intermediary in case of any potential disputes, b) these higher transaction costs practically eliminate the possibility of small casual transactions, c) merchants

[1] https://news.gallup.com/poll/110914/majority-americans-angry-about-financial-crisis.aspx Accessed on 12-June-2018

are wary of the credibility of online payments made by their customers, and hence, customers need to give up more information about their identity than would be necessary for a direct cash transaction, and d) a reasonably small proportion of fraud in online transactions in accepted as unavoidable. “These costs and payments uncertainties”, the paper stated, “can be avoided in person by using physical currency, but no mechanism exists to make payments over a communications channel without a trusted party.”

It is plausible that Satoshi Nakamoto – the pseudonymous author of that document – invented the bitcoin cash system to reduce the reliance on financial institutions as trusted intermediaries for enabling online transactions between two or more parties. If one were to sum up the core value proposition of the bitcoin cash system then it can be said to be: transactions without intermediaries. It is not hard to see why such a value proposition gained significant traction in the immediate aftermath of the global financial crisis of 2008. The underlying technology that enables the bitcoin cash system is called blockchain.

2.2 So, What is Blockchain?

There are many descriptions of what blockchain is. A. Tapscott and D. Tapscott (2017) describe blockchain as a “vast, globally distributed ledger running on millions of devices, it is capable of recording anything of value.” Swan (2015) compares it to “a giant interactive spreadsheet that everyone has access to and updates and confirms that the digital transactions transferring funds are unique.” Some preliminary thinking at the Federal Reserve describes blockchain as one “specific type of distributed ledger . . . which adds changes to the database via a series of blocks of transactional data that are chronologically and cryptographically linked to one another” (Mills et al., 2016). It appears that popular descriptions of blockchain are converging toward a notion of a distributed database or a ”distributed ledger technology” (Walport, 2016).

While these descriptions are technologically accurate they may not be comprehensive enough to fully capture the general phenomenon of blockchain. Moreover, framing blockchain as a distributed ledger may not attract the kind of attention that this technology deserves from policymakers due to its potential policy implications (and applications). At best it may trigger a couple of questions: a) if it is a distributed ledger that anybody can update (like Wikipedia), then how can we trust its accuracy? and, b) if information about financial transactions is publicly available then what privacy implications does it have? While these are important questions, they do not cover the entire spectrum of potential implications associated with blockchain. Hence, the remainder of this chapter instead of providing yet another description or definition of blockchain, attempts to describe the basic phenomenon of blockchain that is relevant for a policy audience and to general public. It is perhaps most useful then to begin with exploring what blockchain *does*, rather than what blockchain *is*.

A careful observer would note that bitcoin's blockchain does three main things:

1.	Creates money	(What a central bank does)
2.	Enables transactions between entities	(What other banks do)
3.	Records information securely	(What ledgers and databases do)

Blockchain creates money in the form of cryptocurrencies. Cryptocurrency is essentially a digital currency that has some cryptographic properties. These properties protect the units of that currency from counterfeiting. They also provide pseudo-anonymity to the users. Bitcoin is the most famous example of a cryptocurrency.[2] That said, not all instances of blockchain create cryptocurrency. R3's Corda and Linux Foundation's Hyperledger Fabric are two popular versions of blockchain

[2] At the time of writing this chapter

focused at serving the financial and other industries but do not create any cryptocurrency in the process. However, these blockchains still produce value of some form for their participants and users. Thus, it may be more inclusive to state the blockchains produce value instead of just focusing of creation of money.

Secondly, blockchain enables direct transactions between two or more entities that are willing to deal in cryptocurrency (and its derivatives). It is widely believed by the blockchain community that no trusted intermediaries like banks, brokers, or dealers are needed to enable transactions. Until blockchain was invented, entities could transact with each other directly only if they transacted in hard cash; for any online transactions a trusted intermediary was necessary. Disintermediating transactions between entities is one of the novel contributions of blockchain. Note that "transactions" in this context is not necessarily confined to financial transactions; it can refer to any exchange of value. For example, a prototype of MIT's MedRec system demonstrates sharing of electronic health records and other forms of medical research information among healthcare providers over a popular blockchain called Ethereum (Ekblaw et al., 2016). Ethereum is also known for being the first blockchain to provide the functionality of smart contracts. Smart contracts are transactions that are automatically executed when the stipulated conditions of the contract are satisfied. In this sense, smart contracts can be construed as automated transactions between entities that are executed at an appropriate time.

While it may appear that there is no need for any intermediary to enable transactions, the reality is a little more nuanced. Any transaction carried out over a blockchain needs to verified and validated by other participants on (and in some cases off) the blockchain. Until consensus on the validity of those transactions is achieved, the transactions are not confirmed. However, this does not imply that there is one intermediary who has the power of holding up a transaction for whatever legitimate reasons they may have. These intermediaries (in some cases called miners) have an incentive structure that motivates them to confirm validated transactions. This incentive structure keeps the intermediaries engaged in confirming transactions while not giving them power to stop or temporarily halt a transaction from execution.

Lastly, blockchain maintains a history of all transactions that it enables along with information about ownership of assets i.e. who owns what. In traditional banking system, such information is typically stored using a combination of a database and a ledger. Due to this similarity in storing information blockchain is sometimes described as distributed ledger technology. For the sake of completeness, it is useful to note that there are some versions of blockchain that store only final ownership information but not the entire history of transactions. Ripple is one example of such a blockchain. After a transaction is processed to update the change in account balance, the transaction is discarded by Ripple's system. It should come as no surprise that a system that produces value by stringing together diverse participants in an incentive structure would need a way to maintain the ownership of the value produced by the participants. Hence, storing ownership (and related) information in an immutable way is one of the basic capabilities required for any blockchain.

A more generic view of what a blockchain does follows from the above discussion. It appears that blockchains do three main things albeit in different ways:

1. Produce value
2. Enable direct exchange of value between entities
3. Securely tracks information about the ownership of value

Now that we have more clarity on what blockchain does, let us explore how does it work.

2.3 How does Blockchain Work?

Since there are at least 792 versions of blockchain (at the time of this writing), explaining how blockchain works necessitates a more generic approach. Hence, this section will first explain a simple use-case using the most popular version of blockchain – the bitcoin-blockchain. Next, a generic framework that abstracts away the details of how any particular version of blockchain works will be presented. The generic framework will also illuminate some of important design choices that are available to developers of blockchain. The discussion on design choices will show why there are so many different versions of blockchain available today.

2.3.1 *Bitcoin's Blockchain Use-Case: Alice Pays Bob*

Let's say Alice bought some goods from Bob and they are transacting using bitcoins (BTC or Ƀ).[3] After buying the goods Alice pays, say, Ƀ2 using her smartphone. Since blockchain uses distributed ledgers to track who owns what, every copy of the ledger has to record the transaction that shows Alice paid Bob. Hence, a message "Alice paid Bob Ƀ2"[4] is broadcast to all "miners". Miners are participants who maintain and update distributed ledgers either out of goodwill or in pursuit of a monetary reward.

A miner who successfully validates (i.e. does the sender have enough funds?) and confirms (i.e. updates the ledger) a set of transactions has a chance at collecting a monetary reward that includes i) a certain number of newly created bitcoins (Ƀ12.5, at the time of this writing),[5] and ii) all the transaction fees (in Ƀ) included by the

[3] Since an accurate description of the inner workings of blockchain can get quite technical, technical-accuracy has been traded-off in favor of easier readability in explaining this use-case.

[4] Of course, actual names are not included to protect privacy; cryptographic identifiers are used instead. A timestamp is also included.

[5] This reward gets halved in amount every 4 years

senders. However, to collect that monetary reward a miner must find a valid solution to a mathematical problem of hashing before someone else does.

What is hashing? Hashing is a one-way mathematical function that takes an input and produces a random but unique set of characters and numbers as output. For example, hashing the word “blockchain” using the SHA256 algorithm produces a random but unique string shown below:

```
> hash.sha256(‘blockchain’)
> EF7797E13D3A75526946A3BCF00DAEC9FC9C9C4D51DDC7CC5DF888F74DD434D1
```

Since it is a one-way function, we cannot find what the input was by analyzing the output. Secondly, any small modification of input to the hash-function produces a wildly different output. In the case below, the input was modified from “blockchain” to “Blockchain” and the output obtained is substantially different.

```
> hash.sha256(‘Blockchain’)
> 625DA44E4EAF58D61CF048D168AA6F5E492DEA166D8BB54EC06C30DE07DB57E1
```

No modifications in the input produce no changes in the output.

In this use-case involving just Alice and Bob, there is only one transaction that needs to be validated and confirmed: “Alice paid Bob Ƀ2”. Now, the problem that miners have to solve is finding a golden number (or nonce) such that when it is hashed in conjunction with the transaction, it produces an output that begins with a predefined number of 0’s.[6] For example, say, the predefined number of 0’s is five. Miners will then have to append a golden number to the transaction and hash it. The miner that finds a number that produces a hash beginning with five 0’s is the one that will receive the monetary reward.

[6] The number of 0’s is predefined by an algorithm based on the number of participating miners. If the number of participating miners increases (or decreases), the number of predefined 0’s increases (or decreases). The predefined number is also called “difficulty of mining.”

The process of finding the solution would look something like this . . .

```
> hash.sha256(‘Alice paid Bob 2 BTC + nonce = 1’)
> 7f1242c49e8263b9addc8bc078efc60b5a4699abc8817dfc602a26a16c48f04c


> hash.sha256(‘Alice paid Bob 2 BTC + nonce = 2’) # nonce is updated
> 0932a13e04d2a2b29efdc33e64b420bdb52e5b4cfa9cdb38e24aa216e0f79778


> hash.sha256(‘Alice paid Bob 2 BTC + nonce = 3’) # nonce is updated
> 75dc43178c137acb01ee05f7ad6ac58bfdebaa01798d09e9b1440f7e93e53eda

...

...

> hash.sha256(‘Alice paid Bob 2 BTC + nonce = 1195103’)
> 0000057d19db021890555d7b85ac14896c916c2d2c774a146c2e04a6f8b4f593
# Found the golden nonce: 1195103,
 (since the hash begins with five 0’s)
```

The successful miner broadcasts the block of confirmed transactions along with the solution to the mathematical problem (i.e. the golden nonce) to other miners. Every miner verifies whether the found solution is correct by re-calculating the hash on the received block of confirmed transactions along with the received nonce. If verified successfully, the miner appends the new block of transactions to the blockchain. This process is called Proof-of-Work, and it ensures that there is consensus among miners regarding all confirmed transactions. If there was no consensus then all the distributed ledgers would be out of sync with each other.

2.3.2 Generic Framework

If we abstract away the details from the above discussion, one can notice that a few generic steps can describe how every version of blockchain works:

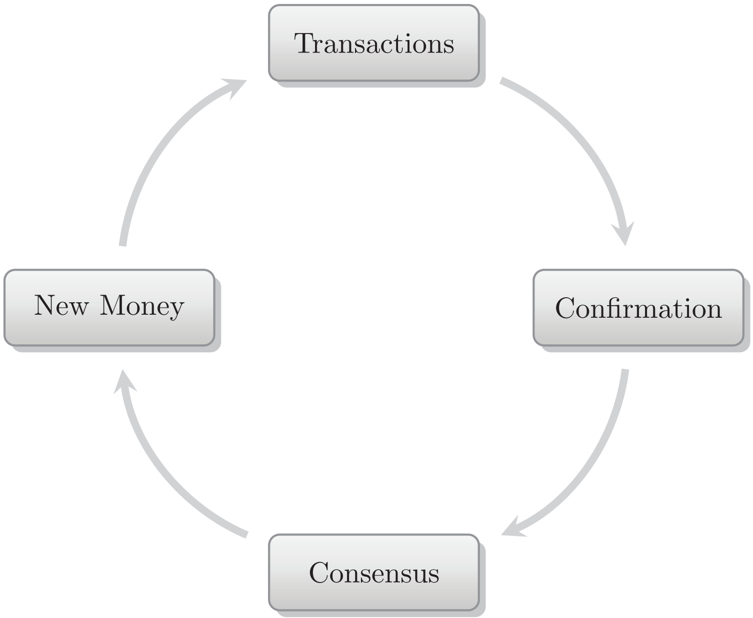

Figure 2.1: Generic Framework: How Blockchain Works

- Exchange of goods, services, or value generates transactions
- Transactions trigger the need for their validation and confirmation
- Confirmation necessitates consensus among miners
- Consensus is achieved by using competition and cooperation

  Competition among miners creates new money into the system

  Cooperation among miners adds a new block of transactions to the blockchain

- New money can enable more transactions

Of course, there are many different ways for implementing each of the above mentioned steps in a blockchain. These design differences are one of the primary reasons why we are witnessing so many different versions of blockchain. The next section discusses some of these major design choices.

2.3.3 Design Choices

Based on a close look at the various versions of blockchain, this section categorizes the design choices made by developers of blockchain into two levels: a) the block level, and b) the platform level.

Block Level

At the block level, it is helpful to think along the lines of the following questions in understanding what factors are governing the design of the blockchain:

- How often is a new block appended to blockchain?
- How is the required consensus for a new block achieved?
- What contents are stored in a block of a blockchain?

These questions not only govern the technical design of blockchain but also its economic design.

Block time appears to be one of the most important design considerations while developing a blockchain. It refers to the average time taken to mine (or append) a new block to the blockchain. It can also be interpreted as how often does a blockchain complete a full circle around the generic framework laid out in Figure 2.1. It is important because it is one of the important determinant of the efficiency of payments on blockchain. Block time is known to vary on average between values as low as 3 seconds (or less) to as high as 10 minutes (or more), at least for the top 10 percent of popular implementations of blockchain. It not only determines how fast does a blockchain grow but also, by extension, determines the confirmation times of transactions. More importantly, it determines the rate of supply of money (or units of cryptocurrency specific to that blockchain). For example, in the bitcoin

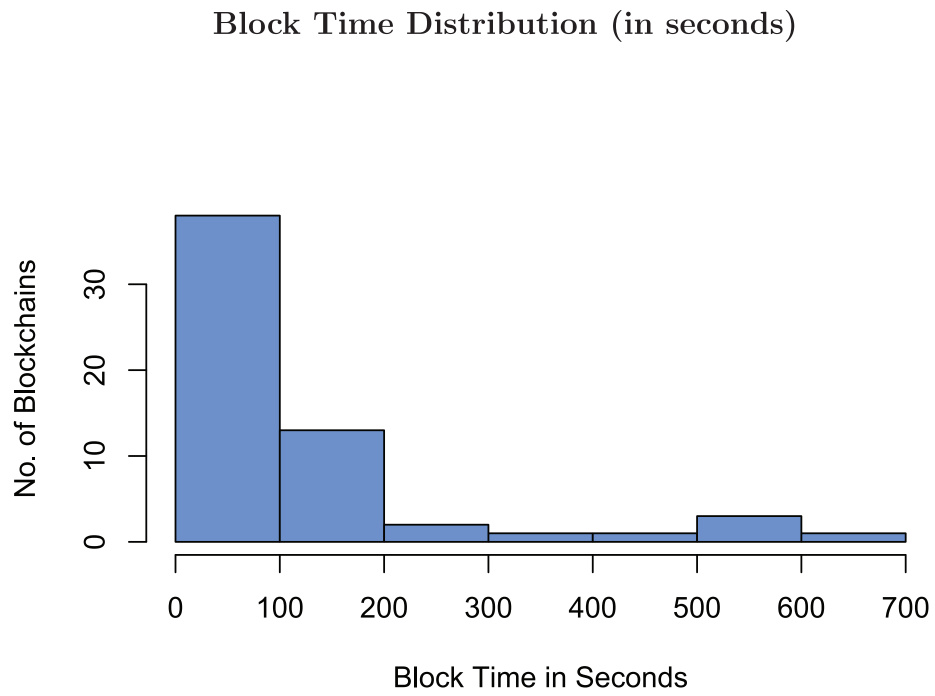

Figure 2.2: Distribution of Block Times of ∼Top 10 percent of Popular Blockchains

blockchain roughly every 10 minutes a few new units of bitcoin come into circulation out of thin air.

The number of new bitcoins that come into circulation also follows a particular algorithm: roughly every four years (or 210,000 new blocks) the number of new bitcoins that come into circulation reduces by 50%. For example, ₿50 would come into circulation every 10 minutes in January 2009 (the beginning of blockchain); it reduced to ₿25 toward the end of 2012, and in 2018 it is ₿12.5. Such an algorithm was chosen apparently because it closely models the rate at which gold is mined.[7] This choice of decreasing money supply results in an upper limit of 21 million bitcoins. Approximately by the year 2140, the issuance of new bitcoins will cease. Broadly speaking, how often new bitcoins are issued and how many of them are issued constitute the 'algorithmic monetary policy' of the bitcoin-blockchain.

Other implementations of blockchain have different block times, and different

[7] See https://en.bitcoin.it/wiki/Controlled_supply Accessed June, 2018

money supply algorithms. For example, Ethereum (another popular blockchain) has a block time of approximately 15 seconds and 3 units of the associated cryptocurrency ether are issued with every new block.[8] Another important difference with respect to bitcoin-blockchain is that Ethereum does not have a pre-determined upper bound on the number of ethers that can be issued. That said, it caps the number of ethers issued per year to 18 million ethers. Another important difference is that Ethereum destroys or burns certain amount of ether with every transaction based on some of the parameters of those transactions. This continuous creation and destruction of ether is expected to reach an equilibrium over time. Thus, even though there is no fixed upper bound on the number ethers that will be issued, the algorithm will converge to a certain interval of ethers over time.

New units of cryptocurrency that are issued (also called as issuance) get assigned to the successful miner of the new block as one part of the block reward. A block reward is provided as a monetary incentive for miners to continue contributing computational power required for the proper operation of a blockchain. Another part of the block reward is the total transaction fees that entities include when they transact with each other on blockchain. The higher a transaction fee an entity includes, the faster the transaction is likely to get confirmed. The transaction fee is calculated differently on different implementations of blockchain. For example, on the bitcoin-blockchain transaction fee is calculated based on the byte-size of the transaction i.e. a transaction that has a size of, say, 512 bytes will be charged more than a transaction that has a size of 256 bytes. However, on Ethereum transaction fee is not calculated by the byte-size of the transaction. Instead, the fee is calculated based on the number of computational operations that are needed to carry out a certain transaction. The more complicated a transaction is, the more fee is charged to process the transaction.

[8] https://www.ethereum.org/ether Accessed July, 2018

$$Block\ Reward = Issuance + Transaction\ Fee$$

Sometimes, in the case of a blockchain with smaller block times, there is a non-trivial chance that more than one miner may have successfully produced a new valid block. In such cases, an additional small block reward is given to those miners. For example, in the case of Ethereum such miners may get 0.625-2.625 ethers for their successful mining effort. One can safely say that the smaller the block times are, the greater are the chances of producing more than one successful miner per block.

However, on the other hand larger block times are associated with a smaller scale (or throughput) of transactions on the blockchain. Roughly 7 transactions are processed per second on the bitcoin-blockchain, and around 15 transactions per second (TPS) on Ethereum. For comparison, VISA servers are known to handle around 24,000 TPS.[9] Blockchain developers are constantly searching for new approaches to increase the scale of transactions on blockchain. Each of these approaches are essentially newer design choices that have either resulted in a new blockchain or in a substantial change in the existing implementation of a blockchain.

Developers have tried increasing the size-limit or block size of a typical block on a blockchain. This approach has the risk of centralizing mining activities because increasing the block size, for obvious technical reasons, necessitates more powerful computational infrastructure to sustain competitive mining. It can lead to centralization of mining activities because more miners would find it more profitable to pool their resources rather than mine the blocks themselves. Over time this may result in only a few miners who can afford to spend on infrastructure are the ones collecting the block reward. Secondly, developers tried decreasing the block size of a typical block *and* the block time (for example, Ethereum). It appears that this

[9] See https://usa.visa.com/run-your-business/small-business-tools/retail.html Accessed July, 2018

approach found larger acceptance. However, it distorts the monetary incentives of miners. A miner who may have been the first to successfully mine a block may not be the one to collect the block reward because of network latency (or delay in transmitting information on the Internet) and various other external factors. If this happens repeatedly the miner may not be able to continue mining activities because the expected block reward may no longer meaningfully outweigh the expected costs of mining.

$$Scale\,of\,Transactions \propto \frac{Transactions\,per\,Block}{Time\,per\,Block} \propto \frac{Block\,Size}{Block\,Time}$$

Another reasonably popular design choice that developers of blockchain are experimenting with is the development of various new consensus algorithms. Consensus, as explained earlier, is needed to grow the blockchain one confirmed block of transactions at a time. Proof-of-Work was first discussed in (Nakamoto, 2008), and it still remains the most popular type of consensus algorithm. As discussed earlier, this algorithm takes a set of unconfirmed but valid transactions and groups them into a block that is then fed as an input to the hash function to obtain a hash value with a predefined number of 0's. The predefined number of 0's or the difficulty level of the blockchain is adjusted based on the number of miners in the blockchain such that the block time on average remains around a desired interval of time (10 minutes for the bitcoin-blockchain). To address issues of scaling due to high block times (and the energy consumption issue), Proof-of-Stake algorithm was developed (King and Nadal, 2012).

In the proof-of-stake algorithm, miners again bundle unconfirmed (but valid) transactions into a block but instead of hashing it, they bet on the validity of that block. When other miners receive a block of transactions, they verify the validity of transactions within that block. If invalid transactions are found, the original miners who were betting on the validity of the block lose their stake. If the block is verified to be valid then all other miners accept the block and the original miners proportionally divide the block reward among themselves according to their stake. This algorithm is energy-prudent and reasonably fast in producing blocks. However, it raises some distributional issues – the rich (who can stake more) get richer, and the poor (who cannot stake more) remain poor. A blockchain that uses such a consensus algorithm may over time witness that only a few miners may be the ones collecting a lion's share of block rewards.

To alleviate the distributional issues of proof-of-stake, an algorithm called delegated proof-of-stake was developed (Larimer, 2014). Under this scheme, participants routinely elect miners or block producers and only these elected block producers have the right to append new blocks of valid transactions to the blockchain. If a block producer acts maliciously not only they will lose their stake but also the people who voted for them lose their stake. On the other hand, after a block producer successfully appends a new block the associated block reward is proportionally distributed among all voters.

In some cases like the EOS blockchain, 21 block producers receive a fixed salary for their services in addition to the variable block reward. While such a consensus algorithm drastically reduces the block time (roughly 3 seconds per block in EOS), it is more centralized compared to other implementations of blockchain. Because it is more centralized and a small number of block producers have more power, EOS community has developed a constitution to govern the behaviors of block producers, voters, users, and developers of EOS. Small variations to these consensus algorithms have resulted in new implementations of blockchain.

There are many other consensus algorithms that are being developed like proof of weight, proof of authority, proof of space-time, byzantine fault tolerance, federated byzantine agreement, and others.[10] Since we are trying to investigate the main

[10] See https://hackernoon.com/a-hitchhikers-guide-to-consensus-algorithms-d81aae3eb0e3 Accessed July, 2018

design choices of blockchain, and consensus algorithm is one among them, it may overwhelm the reader to cover all these algorithms in this section. Any reader interested in knowing more about consensus algorithms may find it useful to refer to (Sankar, Sindhu, and Sethumadhavan, 2017), (Mingxiao et al., 2017), and (Nofer et al., 2017) for more insights on consensus algorithms.

In addition to consensus algorithms, block times, and block size what contents are stored in a typical block of a blockchain is also an important design choice. The block-contents design choice is expressly focused on increasing the privacy of transactions on a blockchain. On the bitcoin-blockchain, even though the users are pseudonymous various techniques can be used for probabilistic identification of a user. Moreover, the bitcoin-blockchain does neither hides the balances held by a pseudonymous id nor does it hide the number of bitcoins transacted between two ids. All information about transactions required for probabilistic identification of users is available in the blocks of a blockchain i.e. the history of all bitcoin transactions is stored on the blockchain. Since blockchains operate in an append-only manner, once the information about a transaction is recorded on the blockchain it cannot be erased. This raises some issues related to privacy and makes bitcoin not easily fungible – a currency is said to be fungible if there is no identifiable difference when one unit of the currency is replaced by another unit of the same denomination. If a bitcoin is used in an illicit transaction then the bitcoin remains forever tainted and users, understandably, may not want to accept it to avoid issues with law enforcement.

Hence, a few implementations of blockchain chose to focus on addressing these privacy and fungibility issues associated with the original implementation of blockchain. Some of these popular implementations are: zerocash-blockhain (Sasson et al., 2014), zcash-blockchain (Hopwood et al., 2016), monero-blockchain (Noether, 2015), and dash-blockchain (Duffield and Diaz, 2014). These blockchains provide the ability to the users to have confidential or shielded transactions, which means information about the amount transacted and the parties who transacted is stored in an encrypted format in a block. This makes it significantly difficult to analyze the transactions for probabilistic identification of users. Innovative cryptographic algorithms like zero-knowledge proofs (Ben-Sasson et al., 2013) are used to carry out the transactions without revealing information to non-parties.

Other implementations chose to not store the transactions on the blockchain to increase the performance of blockchain or the speed of confirming transactions. For example, blockchains like Ripple, Stellar and others that work with mainstream banking and other financial institutions store just the final account balances in their distributed ledgers. Transactions are discarded once the corresponding balances are updated in those ledgers. This makes the process of achieving consensus much faster compared to other consensus algorithms. For example, Ripple's blockchain is known to support 1500 transactions per second.[11]

Platform Level

The first implementation of blockchain created money in the form of bitcoin. It did not have any other functionality apart from creating money and enabling transactions between parties. However, that changed with the development of Ethereum in 2014.[12] It upgraded blockchain's functionality by providing the ability to execute computer code on the blockchain. This ability to program on the blockchain enabled a platform that: a) supports the development and deployment of distributed applications (or software that run on multiple systems at the same time), b) enabled the creation of smart contracts or automated payments that trigger when certain stipulated conditions are met, and c) enables the creation of tradeable digital as-

[11] See https://ripple.com/xrp/ Accessed July, 2018

[12] See https://www.ethereum.org/ Accessed July, 2018

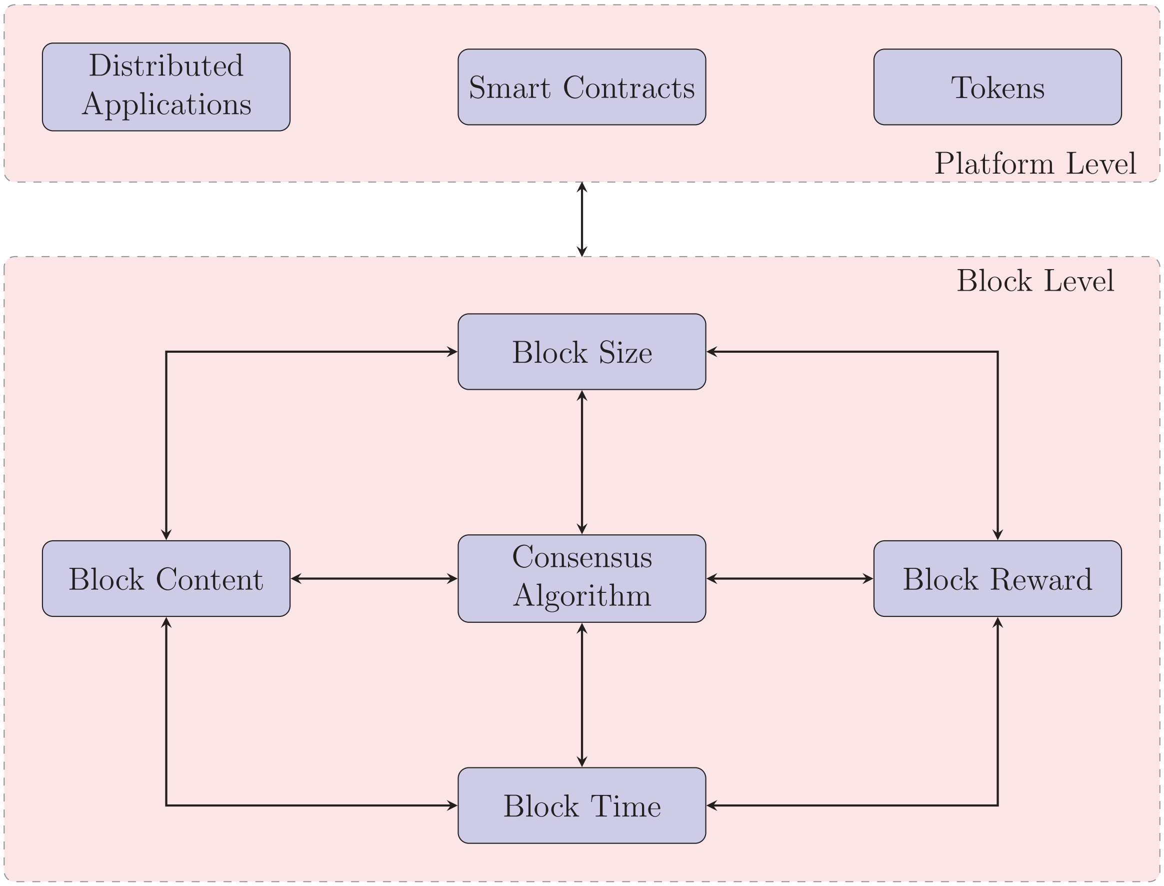

Figure 2.3: Design Choices and their Relationships with Each Other

sets called tokens. Many different blockchains with associated platforms have been created since then for example, NEO, EOS, Cardano, and QTUM.

These abilities to create distributed applications, smart contracts, and tokens have contributed to a lot of new developments on blockchain. Two developments that are worth mentioning here are new avenues for crowdfunding in through initial coin offering (ICO), and distributed autonomous organizations. Anyone with a reasonable new idea for a distributed application can raise funds for its development through an ICO. Investors can buy the tokens being sold through the ICO and either hold them or trade them. Owning these tokens is similar to owning shares of a company. The developers of the application can use the funds raised through the ICO to develop the application. Once the application is deployed, based on its usage the value of the tokens may rise (or fall).

Distributed autonomous organization (DAO), are organizations that no one entity owns. It is an organization that resides on the blockchain in the form of code. The bylaws governing the actions of this organization are implemented through smart contracts. Shareholders vote on the precise functions that a DAO is required to fulfill. It operates in a way that maximizes the value of its shareholders by providing services in the free market.

One obvious fact that needs to be mentioned for the sake of completeness is that all these design choices (at the block and at the platform level) are made while giving substantial importance to the security of blockchain. It goes without saying that without a reasonable level of security a blockchain cannot sustain for long. With this understanding of what does a blockchain do, and how does it work, we can proceed to holistically look at the phenomenon of blockchain.

2.3.4 The Phenomenon of Blockchain

The root meaning of the word phenomenon comes from the Greek word *phainein*, which means 'to show' or *phainomenon* to show the object that is appearing to one's view. This section will therefore attempt to show a holistic view of blockchain that does not rely heavily on the technical details of blockchain.

As discussed in the previous sections, the first blockchain was invented in the immediate aftermath of the global financial crisis of 2008. The core value proposition of the first blockchain can be summed up as: financial transactions without financial intermediaries. Since there was a general atmosphere of anger toward financial intermediaries like commercial banks, investment banks, dealers, brokers, and insurance providers in 2008-09, blockchain's value proposition found some traction with a small subset of the general public. To actualize its value proposition, blockchain's developers first invented a digital currency called bitcoin that has some cryptographic properties that prevent duplication of its ownership. Secondly, they used bitcoin as a monetary incentive for miners to carry out some of the tasks of a traditional intermediary like verifying, validating, and confirming transactions but the miners do not have any legitimate power to unilaterally stop any transaction. Following the relative success (and some severe limitations) of the bitcoin-blockchain, hundreds of new blockchains have mushroomed over the Internet in the past decade. Each blockchain either differs in the way it is implemented, the incentives it provides, the participants it allows, or the value it produces.

Nevertheless, there appears to be a central theme that can unite all these blockchains into one category. Every blockchain directly brings together two or more parties interested in exchanging value. It also has an incentive structure for irreversibly confirming the exchange of that value. Lastly, after transactions of value are confirmed by the participating entities, new units of cryptographically protected value are provided as incentive. Based on this central theme, one can define blockchain as *a technology that enables its users to create and exchange cryptographically protected value with each other.*

One can notice that the act of bringing entities together to enable them to exchange value is generally referred to as “market-making”. Traditionally, we have witnessed two types market-making entities in the market: a) brokers who bring buyers and sellers together and charge a commission for it, and b) dealers who act as buyers when there are no buyers and act as sellers when there are no sellers and in that process profit from the bid-ask spread (or by buying low and selling high). Dealers, thus, are generally known to undertake higher risks on behalf of the transacting entities compared to that of brokers. The two important kinds of risks market-makers undertake are (of course, they take other kinds of risks as well): a) settlement risk – the risk that the payment may come in late or may not come in at all, and b) counter-party risk – the risk that once the payment is made the asset or good is not delivered on time or not delivered at all.

Blockchain, however, appears to be changing the dynamics of market-making in some subtle ways. First, in the case of markets enabled by blockchain it appears as if the market-maker is the market (or vice versa). This is because the buyer and seller, if they so wish to, can directly transact with each other (at least for purely digital goods or assets) on a blockchain. Of course, whether people are approaching blockchain as a market in itself (or just as a payment system) is a different question that will be explored later.

Secondly, since blockchain does not undertake any settlement or counter-party risks in making the market, the transacting parties are responsible for managing their own risks. Again, questioning whether shifting risks to transacting parties is a good thing is a different matter that we will go into later.

Lastly, when transactions are irreversibly confirmed many blockchains create new units of value in the system. This seems to be the most peculiar characteristic of blockchain. In many blockchains, new units of cryptocurrency get generated after a set of transactions that have occurred have been irreversibly confirmed. Over the past few decades, new units of fiat currency have been created only at the discretion of a central entity like the Federal Reserve that implicitly takes into consideration many other factors than just the fact that transactions have occurred.

Based on this, one can safely say that *the phenomenon of blockchain appears to be about changing the dynamics of market making.* While it may strike as an interesting phenomenon to observe, it is not clear yet what it may imply for markets, for intermediaries, for policymakers, and for the public in general.

2.3.5 Next Steps

Hence, the next few chapters of this dissertation will further examine this phenomenon with the intention of identifying what can be some of the main policy implications of blockchain. Also, with the intention of assessing whether policymakers should support blockchain, oppose it, closely watch it, or ignore it altogether. We will first examine whether blockchain is eliminating intermediaries or just replacing one kind of intermediary with another kind. If it is truly eliminating the need for intermediaries, should policymakers consider regulating blockchain and if so how can they go about it? If it is just bringing newer intermediaries in place of traditional intermediaries, then can the existing policies be applied to these new intermediaries with minor modifications if necessary? Or is a fundamentally different set of policies needed, if at all?

3. Intermediaries

> "*It is well-known what a middleman is; he is a man who bamboozles one party and plunders the other, . . .*"
>
> —Benjamin Disraeli

If the above characterization of middlemen by Benjamin Disraeli[1] (former Prime Minister of United Kingdom) from 1858 feels right or is relatable, a question that may arise is: how is it that middlemen have survived to this day? One plausible answer is that middlemen survived because the services they provide are worth more than the costs imposed by their bamboozling of one party and plundering of the other. Another plausible answer is that there was no alternative to middlemen, at least, until now. As discussed in the previous chapter, one of the core value propositions of blockchain is to enable market transactions without intermediaries. After witnessing a decade of existence of blockchain and hundreds of its variants being developed and deployed, it may be useful to explore whether blockchain is truly eliminating the need for intermediaries that enable market transactions.

3.1 Is Blockchain Eliminating Intermediaries?

A meta answer to the above question is that blockchain has not been able to completely eliminate the need for intermediaries. For example, the need for an escrow[2] agent to temporarily hold the payment in bitcoins until the goods are delivered to

[1] See The Saturday Review, July 24, 1858, page 77

[2] An "Escrow is a legal concept in which a financial instrument or an asset is held by a third party on behalf of two other parties that are in the process of completing a transaction." See https://www.investopedia.com/terms/e/escrow.asp

the buyer shows that the need for intermediary continues to exist for bitcoin transactions. Of course, the specific relationship that an escrow agent has with buyers is slightly different than a bank has with the buyers when they buy goods online using debit card, credit card, or through Internet banking. The escrow agent does not necessarily bring the buyer and the seller together for a transaction, but merely completes the payment for that transaction when the stipulated conditions are met, like the goods have been delivered to the buyer.

Put simply, escrow agents safeguard buyers from counterparty risks and sellers from settlement risks, and they usually charge a small fee in return. It should not come as a surprise that there is an implementation of a blockchain whose main aim is to provide trust-worthy escrow services for cryptocurrency transactions. The native cryptocurrency of this blockchain is called EscrowCoin.[3] On this blockchain, entities that provide escrow services for transacting parties get rewarded in EscrowCoins just like miners get rewarded in bitcoins for validating and confirming transactions on the bitcoin-blockchain. Regardless of there being a blockchain for escrow services, the need for a trusted third-party who can act as an escrow-intermediary continues to exist.

On the other hand, developers of the Ethereum blockchain cleverly automated the role of an escrow agent using smart contracts to remove the need for an intermediary to temporarily hold payments. Smart contracts hold the payment from the buyer until the seller delivers the specified goods to the buyer. After the smart contract receives valid information that the goods have been delivered to the buyer (and that the other stipulated conditions of the transaction have been met), the payment is released to the seller.

However, the problem lies in the fact that Ethereum cannot access information that is not already available on the blockchain. To get the information about whether

[3] See https://escrow-coin.com/ Accessed in Dec-2018

specified goods were delivered to the buyer, the smart contract has to rely on a trusted third-party to provide that information. Such a third-party is popularly called as an oracle. The main role of an oracle is to provide all the necessary off-chain information or information that is not available on the blockchain for the smart contracts to execute properly. Of course, the oracle charges a small fee for making that information available for that smart contract.

Another kind of intermediary whose significance often gets overlooked in the discussions of the role of intermediaries in blockchain enabled transactions, is the intermediary who delivers the actual goods like FedEx, UPS, or United States Postal Service (USPS). The intermediary who delivers the goods to the buyer is an essential entity in completing a transaction. Blockchain has not been able to eliminate the need for such an intermediary.

Lastly, there is another entity that in a way acts like an intermediary but for some reason is not explicitly considered so in the discussions involving intermediaries and blockchains. These are cryptocurrency exchanges that provide an app called 'blockchain wallet' that the users use to store their cryptocurrencies. A crypto-exchange not just exchanges cryptocurrency for fiat currencies or other cryptocurrencies but also stores the cryptocurrencies of users on its online-servers. Cryto-exchanges have abstracted out a lot of technical complications of holding, using, and exchanging cryptocurrencies through their wallet app. This also means that if a crypto-exchange wants to stop its users from transacting cryptocurrencies (for whatever reasons), it has the ability to do so.[4] Basically, a crypto-exchange manages all the cryptocurrencies that a user holds through its wallet app, just like a bank manages fiat currency for an account holder.

To summarize, there appear to be at least 4 different kinds of intermediaries op-

[4] Binance, a prominent crypto-exchange, had suspended its services for its users in mid 2018. See https://finance.yahoo.com/news/binance-halts-trading-over-abnormal-051504282.html Accessed Dec-2018

erating in the blockchain space. First, a ‘escrow-intermediary’ provides escrow services for blockchain based transactions. Second, an ‘information-intermediary’ like an oracle that provides information for a smart contract to act like an escrow agent. Third, a ‘delivery-intermediary’ like FedEx or USPS that delivers the physical goods or assets to the appropriate transacting party. Lastly, a ’exchange-intermediary’ that abstracts away all the technical complications of handling cryptocurrencies.

The main point that falls out of the above discussion is that even though the core value proposition of popular blockchains appears to be to disintermediate transactions, the various potential risks associated with market transactions necessitate intermediaries. However, toward the end of the previous chapter we noticed that in the case of blockchain the market itself is the market-maker (or vice versa) implying that there was no need for a market-making intermediary. When these two facts are juxtaposed, a potential contradiction arises: blockchain based transactions are intermediated and disintermediated at the same time.

This spurious contradiction gets resolved when we dig a little deeper into what a transaction means. When the seller and the buyer decide to exchange goods for money with each other, a transaction is initiated. The delivery of goods to the buyer along with the final release of payment to the seller marks the successful completion of that transaction. With blockchain, only the *initiation of transactions* is being disintermediated but not the *completion of transactions*. Completion of transactions still depends on intermediaries like escrow-agents, oracles, goods carriers, and crypto-exchanges. Nevertheless, with the advent of blockchain the dynamics between intermediaries and transacting parties appear to be changing.

3.2 Implications of Blockchain based Transactions

If we dig a little deeper into the changing dynamics between intermediaries and transacting parties, we can notice that there are at least two important sources of policy implications. The first source of policy implications comes from illicit trade transactions between buyers and sellers. Second, policy implications due to intermediaries who could, in future, try to improvise on their services as their business strategies evolve. This section discusses these two sources of policy implications in greater detail.

3.2.1 Implications due to Illicit Transactions

Transactions enabled by blockchain and cryptocurrencies may have some important policy implications if the goods being traded are of illicit nature like scheduled drugs, illegal firearms, footage of child pornography, stolen information like credit cards, passwords and personally identifiable information (PII). Since blockchain makes it easy for buyers to directly initiate transactions with sellers of such goods, and since cryptocurrencies make it harder to link transactions to identities of transacting parties it is not unwise to expect that the volume of transactions involving such illicit goods is only likely to increase in the future. Foley, Karlsen, and Putniņš (2018) estimate that close to $76 billion of illegal activities like prostitution and drugs per year are enabled by bitcoin. They also note that bitcoin's share in enabling illicit trade is declining as more opaque alternative cryptocurrencies (Monero, for example) are being used to better cover their tracks.

In the case of delivery of physical goods, private carriers like FedEx etc. have the legal basis to open and inspect most international packages. However, the USPS is generally prohibited from opening any international or domestic shipments and mail.[5]

If policy makers desire to reduce the volume of such illicit trade activities there may be value in modifying existing policies or developing new policies that account for blockchain and cryptocurrency enabled transactions. Since the initiation of transactions can be completely disintermediated by blockchain, it will be extremely difficult to enforce the regulation of initiation of such transactions. The smarter way to regulate may be to regulate the intermediaries who operate on the completion-side of those transactions.[6]

If policy makers want to alleviate this issue, the basic puzzle that their policies must aim to solve is the unambiguous linking of transactions to identities and locations of transacting parties. As long as law enforcement officials can get access to information that enables them to link transacting parties to transactions (goes without saying with proper authorizations), the policy problems associated with blockchain enabled transactions of illicit goods can be kept in check. The following table tries to identify the information gaps that may be needed to be filled to address this policy problem. It also shows what information collecting instruments are in place, for example the Know Your Customer (KYC) regulations have been put in place for some entities.

[5] See the Office of Inspector General's report on "Use of Postal Service Network to Facilitate Illicit Drug Distribution" https://www.oversight.gov/sites/default/files/oig-reports/SAT-AR-18-002.pdf Accessed Dec-2018

[6] That said, it may still be beneficial to ban or outlaw the initiation of illicit goods' transaction even if it is not enforceable because it may give alternative avenues for law enforcement officials to bring criminal perpetrators to justice.

	Information About	
	Transacting Parties	*Transacted Goods or Assets*
Escrow-Intermediary	KYC (in place)	[Info Gap]
Information-Intermediary or Oracle	[Info Gap]	[Info Gap]
Delivery-Intermediary	[Info Gap]	Location (in place) Type of goods [Info Gap]
Exchange-Intermediary	KYC (in place)	[Info Gap]

Information Gaps in Linking Blockchain Transactions to Identities (if needed)

However, one must note that the above approach to policy-making has some short-comings. First, it can work only as long as blockchain developers do not succeed in their attempts to disintermediate the completion-side of transactions. Once the process of completing initiated transactions is disintermediated by blockchain (or by other means) new policies may be needed. Whether the process of completing transactions can be disintermediated will remain an open question that only time can answer. Second, it can mainly apply to illicit physical goods but not digital goods. With digital goods like stolen credit card information, PII, or corporate secrets the process of delivering those goods may not necessitate an intermediary at all if transacted using blockchain. Hence, it may be extremely hard to regulate blockchain based transactions involving illicit digital goods as there are no interfaces that can be regulated in such transactions.

3.2.2 Trend of Decentralization

It may be important to explore the trend of decentralization because it may some important implications for law enforcement agencies if the trend grows. Developers of blockchain are continuing their effort to decentralize the completion-side of transactions. However, to what extent they can be successful at this is still unclear because of the various complications involved in implementing it for all cases. The concept of 'decentralization' at least in the blockchain community has been overloaded with meanings. It is one of the reasons why this dissertation has not used that concept until now. Moreover, because the descriptions of decentralization can be confusing we will not even mention those explanations here. Instead, this section aims to explore what decentralization means within the context of this dissertation.

Decentralization means that even though an interim service or task is essential for blockchain enabled transactions to succeed, anyone can opt-in to act as that intermediary for a particular transaction. There is no one central entity responsible for providing that service. Decentralization of intermediate tasks or services is a trend that is continuing to grow in the blockchain space. In fact, all of the previously identified intermediaries are in some way being decentralized. First, (as mentioned earlier) an entire blockchain is being developed to provide escrow services in a decentralized manner. Anyone can run a client-software of this blockchain and start providing escrow services and get paid in EscrowCoin (ESCO) for providing those services. The more number of such escrow services providing intermediaries are on this blockchain the more decentralized it is.

Secondly, automated oracles are aggregating off-chain information from various sources in a decentralized fashion and making that information available on the blockchain. ASTRAEA as conceptualized by Adler et al. (2018) and Augur[7] are a couple of examples of automated oracles. Any entity can opt-in to provide information required for smart contracts to execute properly, and the automated oracle aggregates the values from all sources and provides the aggregated value to the blockchain for proper execution of smart contracts. Some of the fees collected by

[7] See https://www.augur.net/ Accessed Dec-2018

smart contracts are then distributed among all individual information-providers for the oracle.

Thirdly, many blockchain-based delivery startups are working on decentralizing the delivery of physical goods to the buyers from the sellers. Triwer (a Norwegian company), ParcelX (a Singaporean company), VOLT (a South Korean company), and PAKET (an Israeli company) are all trying to decentralize the delivery of goods from source to destination.[8] They are either trying to develop an Uber-style crowd-sourced delivery model or a completely decentralized peer-to-peer physical delivery protocol.

Lastly, the role of crypto-exchanges are also being automated in a decentralized way. There are blockchain-based decentralized exchanges (or dex) like IDEX, Waves Dex etc. that are providing decentralized peer-to-peer cryptocurrency exchange services. There are no central owners of such exchanges. These decentralized exchanges enable traders to trade with each other directly.

As of now, the decentralization of these interim services is still in its nascent stage of development. The underlying inefficiencies of blockchain only render these services even more inefficient when decentralized. Another general problem with such services is that they are hard to use. For example, decentralized exchanges tend to complicate the process of handling cryptocurrencies as compared to the centralized exchanges. Furthermore, the risk has to be borne by the users, which has been an obstacle in the adoption of such decentralized exchanges for cryptocurrency trades. That said, these platforms are continuing to grow and newer ones are coming into this space. When (and if) blockchain developers invent new ways of making their platforms efficient, these decentralized services may see a big jump in their adoption.

That said, incentives provided to attract more participants are showing signs

[8] See https://hackernoon.com/parcel-deliveries-on-the-blockchain-26d8a15f2712 Accessed Dec-2018

of undercutting decentralization of those services. A prime example of this phenomenon is the growing size of mining pools. Many miners are finding it easier to gain reward for their mining services by signing up to a mining pool, where the mining rewards received by the pool is distributed among its members. Ironically, rewards provided as incentives for decentralizing a service seem to be encouraging centralizing behavior among the participants. This is an incentive design problem in blockchains with no known solutions at the time of this writing.

Policies that can regulate centralized entities can easily be instrumented to keep a check on illicit transactions that take place using centralized services. However, decentralized services may pose tougher policy challenges and they may be harder to solve. The main puzzle that policymakers will need to solve, however, will remain the same – linking identities to transactions. The one good thing about decentralized services is that all the information remains publicly accessible to ensure any entity can opt-in to provide those services. If policymakers want to crack down on illicit transactions enabled by blockchain, it may be useful to fund academic and industry research that can use publicly available information to link identities with transactions at least in a probabilistic manner if not in a definitive manner. A probabilistic linking could at least result in probable causes that can enable further investigations.

It is ironic that a technology that was developed to provide more privacy of transactions may actually give impetus to more intrusive policies mainly because the 'bad guys' can (and probably will) exploit that technology. This is a strange interaction of technology and public policy that not just surfaces in the case of blockchain but one that surfaces with the development of most privacy-enhancing technologies. This may suggest that the onus of preventing the criminally or politically motivated actors from exploiting a technology lies as much on the developers of that technology as much as on policymakers and law enforcement officials who want to regulate its use.

3.2.3 Implications due to Evolution of Intermediaries

If crypto-exchanges are holding the cryptocurrencies of their customers, then is it possible for them to act like a bank and lend the deposited cryptocurrencies to others for profit? If an escrow agent is holding cryptocurrencies for a few days until a transaction is completed, can that agent invest that money during the time-interval before releasing the payment to the seller? Can delivery service providers also provide escrow services for their customers? A conditional-answer to these questions is that if those intermediate entities want to broaden their services, they can do so; and if it is profitable to broaden their services they are likely to provide such services. While all of the possibilities of evolution alluded to through the above questions are technically feasible, not all of them have materialized for unknown reasons. Thus, we will focus on the possibility that has veritably materialized: crypto-exchanges are providing borrowing and lending services.

Borrowing and lending services provided by crypto-exchange platforms is a genuinely growing business. There are many different cryptocurrency platforms providing various varieties of borrowing and lending services. Before the invention of blockchain and cryptocurrencies, currency borrowing and lending was limited to the fiat-to-fiat variety. Now, however, there are additional varieties like crypto-to-crypto, fiat-to-crypto, and crypto-to-fiat.

Because the crypto-to-crypto variety of borrowing and lending has the potential to be completely decentralized, a coherent policy framework is still emerging. However, it may be useful for policy makers to note that any regulations aimed at regulating purely crypto-to-crypto exchanges may be extremely difficult to enforce due to the completely decentralizable nature of such exchanges. Hence, most of the immediate policy implications that can be addressed emanate from the crypto-to-fiat and fiat-to-crypto exchanges. These policy implications are, by and large, addressable because wherever there is conversion to fiat involved it is likely the case that there is an aggregating- or a centralizing-entity who can provide information on the borrowers. Regulating that centralizing-entity can indirectly enable policy makers to bring cryptocurrency exchanges within the confines of the existing policy framework.

If the wildly fluctuating exchange rates of cryptocurrencies stabilize, it appears fairly possible that a fractional reserve banking kind of system dealing in cryptocurrencies may emerge. Most of the institutional infrastructure needed for a such a system appears to be falling in place. Hence, if the exchange rates stabilize and the number of transactions that can be processed on blockchain matches up with the existing standards of transactions processing set by commercial entities like Visa/-Mastercard then policy makers may have to gear up to the possibility of a fractional reserve system of banking dealing in cryptocurrencies. The main policy implication perhaps of such a reserve banking system emerging in the cryptocurrency space is that there would then be an alternative to the fiat based fractional reserve system governed by the central bank. While having alternatives may be a good outcome in view of the competition that this possibility brings with it, it may have consequences that are perhaps hard to even speculate at this stage. Hence, it may be worthwhile for policy makers to keep an eye on the developments in this space.

While institutionalized escrow agents may find it difficult to invest and profit off of the crypto-assets they hold on behalf of their clients due to the fiduciary responsibilities, independent and unregistered escrow agents may be able to invest the assets they hold in escrow. Nevertheless, even though this possibility can materialize it may not pose a big policy challenge as the escrow agents that are unregistered are unlikely to be involved by buyers and sellers that are making large transactions with each other because of the obvious risks they may face from such an escrow agent. It may be hard for transacting parties to trust such independent and unregistered escrow agents when transacting large sums of money.

Delivery intermediaries may find it logical and profitable to provide escrow services to transacting parties. If they provide such services, then they will have to follow the regulations laid down for providing escrow services. Hence, one may not foresee too many policy risks emanating from such a development. In fact, encouraging this could make it easier for policy makers to unambiguously link identities of transacting parties with the goods being transacted. Hence, it may be useful to encourage escrow service provision by delivery intermediaries if they want to mitigate the risks of illicit transactions enabled by blockchain and cryptocurrencies.

3.3 Summary

In this chapter we saw that even though blockchain has been widely touted as the killer of intermediaries, the reality is that blockchain has been unable to eliminate the need for intermediaries. This is because even though intermediaries may not be essential anymore for transacting parties to initiate a transaction on blockchain, they (the intermediaries) do absorb some important risks on the completion-side of such transactions. Hence, as long as such risks exist intermediaries will continue to play a role – the roles they play may vary from the traditional roles intermediaries have played nevertheless they are needed.

We also noted that policy makers could mitigate the risks of illicit transactions enabled by blockchain by regulating the intermediaries. The main puzzle that policy makers need to solve to address this policy issue is the unambiguous linking of transactions to the identities of transacting parties. By bringing the intermediaries under the regulatory framework, policy makers could mitigate the risks of blockchain-based illicit transactions.

We also saw, however, that there is an increasing trend of the role of intermediaries being decentralized. This trend of decentralization of intermediate services may be problematic from a policy standpoint as enforcing regulations on highly decentralized networks may prove to be rather difficult.

Lastly, we also discussed if there any foreseeable policy issues that may emerge if the intermediaries decide to provide more services than they are providing already. Such an evolution of the role of intermediaries does not appear to be too problematic except in the case of there emerging a fractional reserve banking system that deals in cryptocurrencies. The policy problems associated with such a system emerging can only be speculative at this stage.

4. Blockchain-Created Value

"*The value of an idea lies in the using of it.*"
—Thomas Alva Edison

If one observes the value created through blockchain, it is easy to notice that it falls into at least two important categories: a) value created by the use of blockchain, and b) value created on blockchain in the form of cryptocurrencies and crypto-tokens. This chapter explores these two categories to identify if they raise any important issues that need to be addressed.

4.1 Value Created by Blockchain

Blockchain as a market-maker creates value by providing an ability for two or more parties to directly transact with each other financially over the Internet without relying on a trusted intermediary. Before blockchain was invented, two or more parties could do that only if they transacted directly in cash. With blockchain there came a peer-to-peer cash system called bitcoin that parties can use to transact with each other directly over the Internet.

Creation of value by blockchain did not directly translate into monetary worth of blockchain because its availability and accessibility is not limited. Because it is an open-source technology, anyone interested could avail it, access it, and use it however they deem fit. Of course, it is notoriously hard to intellectually access the technology for non-tech-savvy users. That limitation was monetized into worth by crypto-exchanges that provided blockchain wallets to make it easier for non-tech-savvy users to use the technology without getting bogged down by its technical details.

There are also commercial versions of this technology that create the same value but through different algorithms. For example, the company called Hedera that has a similar technology as that of blockchain calls its product as Hashgraph. While the technical details of the patented hashgraph technology can convince an informed reader that hashgraph's potential efficiency maybe much better than its open-source equivalents, it is yet to be popularly adopted for those claims to have been stress-tested. Interestingly, Hedera markets itself as the "the trust layer of the Internet."[1] While most open source efforts of blockchain are aiming to provide a "trustless" way of transacting on the Internet, it appears that the closed-source efforts are trying to provide a layer of trust for the Internet, and they are trying to monopolize it through patents.

This highlights an apparent contradiction in the blockchain developing community's concept of trust that may be useful to clarify before moving ahead. Trust is basically the courage to deliberately overlook something that can go wrong. Someone who Murphies a Law out of it, may quip "something that *can* go wrong, *will* go wrong." Such a person is likely to seek a trustless system. Someone given to practicalizing the world out of trust (instead of idealizing the world), may argue that to operate in a "practical" world we need trust. Such a person may support a layer of trust on the Internet.

From the perspective of policy makers, however, the latter option may be preferable over the former because of law enforcement reasons (as discussed in the previous chapter). The most generic type of policy issue that is likely to arise with blockchain's value creation as a market-maker is that, as discussed previously, it may enable markets for illicit goods and services.

[1] See https://www.hedera.com/ Accessed Feb-2019

Blockchain as a secure recording-keeping mechanism is currently being experimented with to assess if it can satisfy some of the needs of supply chain. Needs such as tracking the movement of items not just within the confines of a company but also include the company's various suppliers and vendors. In other words, existing supply chain mechanisms are good at tracking the movement of items only within the boundaries of a company. They cannot track items across their various third-party suppliers and vendors. With blockchain, movement of items can be tracked from their origin to their destination across vendors, suppliers, dealers, and retailers.

For example, Walmart with the help of IBM is currently experimenting the use of blockchain in tracking food items from farm to retail sales counter.[2] The stated motive behind it is to achieve better food safety. How can blockchain help in achieving food safety one may ask. The way blockchain has found usefulness in food safety is that if there is an outbreak of say E. coli because one of the suppliers unknowingly had supplied lettuce infected with the concerned bacteria in them. The supplying entity itself may have procured lettuce from various different farms among which, let's say, lettuce grown on one farm was infected with the bacteria. If all suppliers of Walmart are using Walmart's blockchain to track their agricultural procurement, then Walmart can precisely track ("within minutes" apparently, and "with pin point certainty")[3] only that batch of lettuce that was supplied by the contaminated farm. They can then pull out of circulation only that batch of lettuce that were supplied from that farm (and perhaps those that came in close contact with it). Without blockchain's record-keeping potential, apparently, Walmart would have to pull out of circulation all lettuce from its shelves regardless of which farm or supplier they were procured from.

The above example shows that blockchain has usefulness in bringing efficiency

[2] See https://www.nytimes.com/2018/09/24/business/walmart-blockchain-lettuce.html Accessed Feb-2019

[3] See https://www.youtube.com/watch?v=SV0KXBxSoio Accessed Feb-2019

to tracking items in complex supply chains – even though the language of marketing focuses more on "food safety", which is a byproduct of the usefulness of blockchain in this case.

Another example of blockchain's value creation as a secure record-keeper is in healthcare information sharing. Pilot studies like MedRec carried out by MIT[4] show that blockchain may create value by satisfying the need for sharing patient information between healthcare providers as needed and appropriate. Since sharing patient healthcare information in the US between healthcare providers is a complicated process, it has raised concerns about the impact it may have on quality of healthcare provision. Blockchain is being used as a secure record-keeper that enables patients (and their primary healthcare providers) to selectively share patient's health information with other healthcare providers as needed.

Of course, the generic policy issue raised by blockchain's usefulness as a secure record-keeper is that of privacy. If these pilot projects have the potential for being used widely, policy makers will need to have proper policy guidelines in place for institutions to follow.

From the above discussion, it appears that there are at least two generic forms in which blockchain's usefulness creates value. Blockchain as a market-maker, and blockchain as a secure record-keeper. These use cases have some generic policy implications. First one has policy implications in terms of potential increases in market transactions for illicit goods and services by providing interested entities with new ways of transacting with each other. Second one has policy implications in terms its increased potential for privacy violations. Especially, since blockchain follows the "append only" model, it may have important privacy implications for healthcare information.

[4] See https://medrec.media.mit.edu/ Accessed Feb-2019

4.2 Value Creation on Blockchain

The previous section looked at value creation *by* blockchain. This section will explore value creation *on* blockchain.

4.2.1 Cryptocurrencies

Perhaps the most widely known example of creation of value on blockchain is bitcoin. Bitcoin is the first cryptocurrency that was created on blockchain. A cryptocurrency can be defined as a cryptographically protected store of value. It appears that cryptocurrencies would not be possible without blockchain. Bitcoin is a form of peer-to-peer cash for online transactions. As the past decade has shown, bitcoin is a form of digital money (or cryptographically protected value) because it has served three widely expected functions of money: a) store of value (albeit quite volatile in the case of bitcoin), b) unit of account (even though bitcoin itself has been treated like a commodity many times), and c) a medium of exchange.

In the case of the traditional monetary system, new units of money enter the economy through central bank's lending, investments (or open market operations), or through return on investments in financial instruments created by the central bank. In the case of blockchain, however, new units of money enter the economy as a reward for providing a service that is central to the continued operation of blockchain. New (and valid) transactions when confirmed through consensus create new units of money for circulation in the economy. Of course, various algorithms are used on different blockchains to determine the number of units of money to be created for circulation.

This difference in the manner in which money is brought into circulation between the traditional case and blockchain's case can be a potential source of concern for monetary policy. That is because there is no central entity deliberately considering whether new units of currency should be brought into circulation in the case of cryptocurrencies. On the other hand, in the traditional case the central bank decides to bring more money into circulation (or to remove some money out of circulation) based on a its monetary policy committee's consideration of the underlying political and economic environments. Some blockchains do take the economic environment into consideration but it is done so in quite a formulaic manner. For example, the rate of supply of new bitcoins reduces by 50% every 210,000 blocks (or roughly every four years). Because blockchains do not take political considerations into account at all and they have a very formulaic view of economic considerations, the money created through blockchains has a potential to affect the efficacy of traditional monetary policy.

How can cryptocurrencies affect monetary policy? If we consider cryptocurrencies as money, then it is fairly obvious to notice that when new units of cryptocurrencies come into circulation they are increasing the total currency or monetary base (the base of money from which other forms of money are created) of an economy. Monetary policy tries to regulate, among other things, the monetary base of an economy. Since monetary policy makers do not seem to have any formal control on the supply of cryptocurrencies, it is clear that blockchains have the potential to undermine the efficacy of traditional monetary policy.

However, the total impact cryptocurrencies have on traditional monetary policy is quite limited at the moment because the total money supply of cryptocurrencies is orders of magnitude smaller compared to total fiat money supply. Also, the exact mechanism through which cryptocurrencies can impact traditional monetary policy may be significantly difficult to trace out because of the dispersed nature in which cryptocurrencies come into circulation. Only if the popularity of cryptocurrencies grows substantially high, then it may noticeably affect traditional monetary policy.

Until then it is likely to be only a small concern. That said, it may be important to note that this concern can quickly grow in its importance if the developers of blockchain find a way to rapidly scale-up the processing of blockchain-based transactions that may in turn spur wide-spread adoption of cryptocurrencies.

4.2.2 Central Bank Digital Currencies

Some central banks appear to be thinking of mitigating the problems of non-centrally issued cryptocurrencies by developing what the Bank of International Settlements (BIS) calls it as central bank digital currencies (CBDC).[5] "CBDC is potentially a new form of digital central bank money that can be distinguished from reserves or settlement balances held by commercial banks at central banks" the BIS report says. e-Krona by Sweden and e-Peso by Uruguay seem to be two rather serious experiments of implementing CBDCs.[6]

Upon deeper investigation it appears that CBDCs are essentially central bank issued cryptocurrencies and crypto-tokens that are created by using, what they call, distributed ledger technology (or blockchain). The prevailing logic behind CBDCs appears to be to harness the benefits of cryptocurrencies such as faster settlements of payments, protection of currency units from counterfeiting, and quasi-anonymity in transactions but issued through the central bank so that the risks posed by cryptocurrencies such as potential for money laundering, terrorism funding, and tax evasion can be mitigated because the central bank will have a greater control on those units.

In the context of regulating financial markets in the US, if CBDCs are issued by the Federal Reserve then it can potentially mitigate systemic risk in the financial market. Currently, this issue is being addressed by the Title VIII of the Dodd-Frank

[5] See https://www.bis.org/cpmi/publ/d174.pdf Accessed Mar-2019

[6] See https://www.bis.org/publ/bppdf/bispap101.pdf Accessed Mar-2019

Act.[7] The Title VIII of the Act aims to mitigate the systemic risks created by "financial market utilities (FMUs) designated as systemically important by the Financial Stability Oversight Council." These utilities "provide the infrastructure for transferring, clearing, and settling payments, securities, and other financial transactions among financial institutions or between financial institutions and the system." If any of these utilities cannot settle their transactions, it produces a cascade of unsettled transactions in the financial system that can lead to a financial crisis if there are massive payment-deadlocks in the financial system.

Because CBDCs (along with blockchain technology) can provide faster and firmer settlement of payments, they can potentially mitigate settlement risks not just those emanating from FMU's but also from all market players transacting in CBDCs. This may reduce the need for some rather intrusive inspections of FMU's by the regulatory entities. Also, CBDC's can potentially provide most of the benefits of cryptocurrencies while mitigating their policy risks because they can be controlled by authorized entities.

4.2.3 Bank Issued Crypto-Tokens

In the US, however, banks like J.P. Morgan Chase and Signature Bank are experimenting with creating crypto-tokens on blockchain. The main difference between a cryptocurrency and a crypto-token is that crypto-tokens are backed by monetary value where as cryptocurrencies represent monetary value themselves. Crypto-tokens are popularly known as StableCoins, because their exchange value with the fiat dollar is fixed unlike that of cryptocurrencies like bitcoin. The token created by J.P. Morgan Chase is known as JPM Coin, and that by Signature Bank is called Signet. Every JPM Coin (or Signet) is backed by a US dollar, and its planned us-

[7] For more details on the Title VIII of the Act see:
https://www.federalreserve.gov/paymentsystems/title-viii-dfa.htm Accessed Mar-2019

age is to settle payments between the bank's wholesale payments business clients. According to Signet developers "[t]ransactions made on the Signet Platform settle in real time, are safe and secure, incur no transaction fees, and require a minimum account balance of $250,000."

While JPM Coin is still in the pilot-test mode and has not received all the regulatory approvals, Signet has reportedly begun its operations after receiving full regulatory approval from New York State Department of Financial Services.[8] The approvals needed that Signature Bank create and operate its own private and permissioned implementation of the Ethereum blockchain (Ethereum was Signature Bank's choice), as the public Ethereum blockchain is not compliant with existing regulations such as bank secrecy, anti-money-laundering (AML), and Know Your Customer (KYC).

The main idea behind using blockchain based tokens for settling transactions between a bank's own clients and international subsidiaries is to reduce the movement of cash back and forth between them. Instead, they could transact in crypto-tokens, and finally when they have settled their transactions they can redeem those tokens for US dollars. Such a way of operating reduces liquidity pressures on the bank. It also reduces the demand for wire transfers of funds between cross-border entities, which is relatively cumbersome and time consuming. Lastly, it helps the bank consolidate their money so that they can efficiently use it to earn better profits.

If monetary policy makers take the "velocity of money"[9] into consideration and if more banks start relying on crypto-tokens, then it may have important implications for monetary policy. It may potentially decrease the velocity of fiat money whereas increase the velocity of crypto-tokens. This is something that has not been discussed

[8] See https://www.forbes.com/sites/benjaminpirus/2019/02/22/signature-bank-already-has-hundreds-of-clients-using-private-ethereum-jpm-coin-still-in-testing/#4bc3750a3359 Accessed Mar-2019

[9] According to the Federal Reserve, "The velocity of money is the frequency at which one unit of currency is used to purchase domestically- produced goods and services within a given time period."

widely anywhere either because it is not that important or it has not caught the attention of commentators. Exactly how does increase in the velocity of crypto-tokens will affect monetary policy is unclear, and only time can shed more light on it; but it may be useful to just list out any potential policy issues associated with crypto-tokens.

Another important policy issue that arises in the case of fiat dollar backed crypto-tokens is the task of ensuring every crypto-token issued is indeed backed by a fiat dollar. Ensuring such a thing could become quite challenging if the tokens are issued and transacted on a private, permissioned blockchain. This is because immutable record-keeping capability – one of the main benefits of public blockchain – is not quite immutable when the blockchain is private and permissioned. The effort needed to manipulate data on the private blockchain would be far easier because the number of participants are likely to be substantially less compared to a public blockchain. Controlling 51% of the computational power of a blockchain becomes far more achievable in a private setting than a public setting. Moreover, there are no time-stamps in blockchains because every block itself acts a time-stamp – all transactions contained in block by definition are processed before the transactions contained in the next block are processed regardless of when those transactions actually occurred.[10] Policy makers need to be aware of this issue if (and when) they audit institutions issuing fiat dollar backed crypto-tokens on a private blockchain. Digital forensics can reveal signs of manipulation, but the actual manipulations themselves like which specific entries were added or removed may be hard to identify.

4.2.4 Security Tokens

Another kind of crypto-tokens that are steadily gaining popularity but also keeping policy makers on their toes are called 'security tokens.' These include asset backed

[10] The algorithms ensure that no double-spending takes place.

crypto-tokens like real estate backed crypto-tokens, or gold backed crypto-tokens, or ‘expectation of profit’ backed crypto-tokens that were offered as part of Initial Coin Offerings (ICO). There has been quite a lot of debate on what is a security token. Only recently these debates have subsided – but not completed eliminated – when the US Securities and Exchange Commission (SEC) legally interpreted that a crypto-token is a security (only) if it is an ‘investment contract’.[11]

The US Supreme Court first laid down the criteria to assess whether a business transaction is an investment contract in *SEC v. W.J. Howey Co. (1946)*.[12] The SEC (based on that case’s judgment) interprets an investment contract as “an investment of money in a common enterprise with a reasonable expectation of profits to be derived from the entrepreneurial or managerial efforts of others.” This is popularly called as the Howey test. Put simply, any crypto-token that passes the Howey test is called a security token. The Howey test basically has four components:

- There is an investment of money or of some form of asset other than money
- The investment is made in a common enterprise[13]
- There exists a reasonable expectation of profits from that investment
- The profits are an outcome of efforts of others (not that of investors)

Based on this test many of the ICO tokens have been classified as security tokens since ICOs are a way for startup companies to raise funds through crowdsourcing methods. When a crypto-token is classified as a security token, all the regulations related to KYC and AML apply to the issuer of those tokens. This has been a major source of conflict between regulators and companies that offer crypto-tokens because

[11] See SEC’s investigative report https://www.sec.gov/litigation/investreport/34-81207.pdf Accessed Mar-2019

[12] See the Supreme Court’s judgment here:
https://cdn.loc.gov/service/ll/usrep/usrep328/usrep328293/usrep328293.pdf Accessed Mar-2019

[13] There is no commonly accepted definition of what a common enterprise is. See https://blj.ucdavis.edu/archives/vol-5-no-2/why-the-common-enterprise-test.html Accessed Mar-2019

the burden of complying with AML and KYC regulations falls on the company offering those tokens. Those burdens can be quite substantial, especially, if the company has not yet started earning any revenue. This is the epicenter of most policy issues associated with value created on blockchain, and there are no clear policy solutions.

To better comprehend the complex policy issues associated with security tokens (or value created on blockchain), it may be useful to discuss the role of the SEC interspersed with a details about crypto-tokens. The SEC was created by the US Congress in 1934 as the first regulator of securities markets in the US. The main tasks of SEC, among other minor things, are: a) to protect investors from fraud and other kinds of risks associated with investments, b) ensuring the securities market is functioning in a fair and orderly manner, and c) to facilitate capital formation.

An ICO is a way for a company to issue crypto-tokens to investors that can serve as tokens of ownership in the company for those investors (like shares for example). In some cases, those tokens can also be used as digital coupons of utility that can be redeemed in value later when the holder of the token avails some kind of a product or service from the company (like airline miles for example). In return, the company gets its funding to commence its operations. Some tokens are transferable and others may not be transferable; it depends on the specific conditions of the offer. Since the tokens are limited in number, if the service or the product offered by the company (or by some other party) generates demand for those tokens, then the worth of those tokens can increase, subject to the usefulness of the product of service being offered. This creates a possibility of profit, and if the investor is buying those tokens because of this possibility, then it is an investment with an expectation of profit.

This expectation of profit coupled with the hype associated with cryptocurrencies, ICOs became quite popular during 2016-2018. This popularity fueled a fear of missing out (FOMO) among speculators that prompted significant increase in

ICO investments. The New York Times reported in early 2018 that "[a]bout 890 projects raised over $6 billion last year [(2017)], a 6,000 percent increase over the year before, according to Icodata.io, which tracks the offerings."[14] Of course, the hype was exploited by a lot of fraudulent ICOs on the Internet in the from of Ponzi schemes. Until this point, the ICOs were not termed as securities by the SEC.

Seeing a flurry of fraudulent ICOs and ponzi schemes, the SEC stepped in to protect investors from potential fraud, and interpreted ICOs to be securities and that all offerings must comply with the existing laws on securities. The new term that is now used instead of Initial Coin Offering is Security Token Offering (STO) or smart securities (since they are based on blockchain).

Any entity that offers to sell securities – smart or not – must register with the SEC or qualify for an exemption under the amendments made to the Securities Act.[15] These registration and compliance requirements have been unwelcome by blockchain based startups because they increase the cost of raising funds to kick-start their operations. A company has not yet started earning any revenue, has not even developed a minimum viable product or service may find it substantially hard to put processes in place for KYC and AML regulations because of the costs associated with them.

There have been attempts by companies to bypass these compliance requirements by arguing that they are creating a cryptocurrency not a crypto-token. For example, Kin Foundation created by the social media company called Kik (or Kik Interactive Inc.) has proposed the invention of an unit called Kin that is "designed, marketed, and offered as a currency to be used as a medium of exchange within a new digital economy."[16]

It is important to note that the SEC does not consider cryptocurrencies as se-

[14] See https://www.nytimes.com/2018/02/05/technology/virtual-currency-regulation.html Accessed Mar-2019

[15] See: https://www.sec.gov/smallbusiness/exemptofferings Accessed Mar-2019

[16] See https://www.kin.org/wells_response.pdf Accessed Mar-2019

curities; only crypto-tokens are being considered as securities at the time of this writing. This can be seen in SEC's chariman Jay Clayton's statement[17]: "while there are cryptocurrencies that do not appear to be securities, simply calling something a "currency" or a currency-based product does not mean that it is not a security. Before launching a cryptocurrency or a product with its value tied to one or more cryptocurrencies, its promoters must either (1) be able to demonstrate that the currency or product is not a security or (2) comply with applicable registration and other requirements under our securities laws."[18]

The underlying problem producing the conflict between the above two points of view is the similarity between a crypto-token and a cryptocurrency. A crypto-token can also be used as money because it can be used as a store of value, a unit of account, and a medium of exchange just like a cryptocurrency. The crypto-tokens like the cryptocurrencies cannot be counterfeited either; they have the same kind of cryptographic protections. The only technical difference between them arises from the algorithm used to create them; more specifically, the level at which they are created. The cryptocurrencies that we know of until now have all been created at the block-level i.e. when new blocks of transactions are added to the blockchain, implicitly new units of cryptocurrencies are generated. However, a crypto-token is created at the smart contract level or the application level that needs an underlying blockchain to be operational. Applications run on top of or make use of blockchain. Blockchain does not need distributed applications for its operation, but distributed applications need blockchain for their operation. Crypto-tokens can be created by a distributed application, but cryptocurrency is created by a blockchain.

The fundamental questions that Kin Foundation appears to be raising are: why is their cryptocurrency being interpreted as a security token? If a cryptocurrency is

[17] This statement was not necessarily a direct reply to Kin Foundation, but it is very relevant

[18] See https://www.sec.gov/news/public-statement/statement-clayton-2017-12-11 Accessed Mar-2019

not a security, then why is Kin that was designed, marketed, and offered to be used as a currency within a limited ecosystem classified as a security?

One argument against Kin Foundation's stance on this rather important policy issue, of course, is the argument of 'reasonable expectation of profit' as laid down in the Howey test. It can be argued that Kin was invented for commercial purposes, and investments into its development come with a reasonable expectation of profits. However, Kin Foundation's lawyers argue that "Kin purchasers were led to expect consumptive use, not profits". By consumptive use they mean, the Kin token can be used in a manner similar to airline miles, or credit card points. The accumulated miles or points can be redeemed in value, but they cannot be sold or exchanged for profit. In fact, Kin Foundation's lawyers seem to have a reasonable counterargument – at least on first glance to a careful reader but who is not a lawyer – for every criteria of the Howey test when applied to Kin. They argue that using a diluted version of Howey test and using the test beyond the context in which it was first applied, the SEC appears to be overstepping its mandate, and it is engaging in, what they call, "regulation by enforcement." Their arguments have the potential to make a careful reader feel that the Howey test is not quite sufficient to determine whether a crypto-offering is a security.

On the other hand, Kin Foundation's case is not completely strong either – but the SEC is yet to point these things out. In the initial versions of the white paper that laid down the concept of Kin, the developers themselves envisioned the implementation of Kin to be like other crypto-tokens based on a public blockchain. The paper stated that "Kin will be implemented on the public Ethereum blockchain as an ERC20 token. [. . .] These advanced features and active ecosystem make Ethereum a natural fit for Kin."[19]

This implies that they originally intended to *implement* Kin as a crypto-token

[19] See page 8 https://www.kin.org/static/files/Kin_Whitepaper_V1_English.pdf Accessed Mar-2019

– not as a cryptocurrency (even thought they designed, marketed, and offered it as a cryptocurrency) – on top of an existing public blockchain called Ethereum that has a cryptocurrency of its own called ether (ETH). No blockchain is known to have more than one native cryptocurrencies (yet); it does not seem to make sense either to have more than one native cryptocurrencies for a blockchain. Only later, Kin's implementation plan seems to have changed to creating Kin on a native (but public) blockchain of its own, instead of creating Kin on a blockchain that already has a native cryptocurrency.

Creating a crypto-token on an existing blockchain is, at least relatively speaking, far easier than creating a new cryptocurrency because a new cryptocurrency requires a new blockchain. Of course, the code for creating a blockchain is easily available because code to create many blockchains is open source. However, for a public blockchain to operate securely and its native cryptocurrency to have a decent market worth, participants have to be incentivized to run a node of the newly implemented blockchain.

Persuading participants to provide their computational power to a duplicate (but new) version of a blockchain is harder than persuading them to do that for a blockchain that has a novel feature that was not available in the original version of that blockchain. Otherwise, it is economically logical for a participant to provide their computational power to an existing original version of a blockchain that has already found public acceptance, and whose cryptocurrency has gained some worth. That way, the participants at least stand a decent chance of winning a few units of that native cryptocurrency as a mining reward.

The following crude analogy may help one solidify the understanding of this difference: developing a blockchain with a native cryptocurrency is like developing an operating system, whereas developing a crypto-token is like developing an app. While both can contribute to innovation in their own ways (iOS and Uber for example), it is likely that the effort in or difficulty of developing a new cryptocurrency (or a new operating system) is substantially higher than that of developing a new crypto-token (or a new app). When creating a new cryptocurrency, it is rather important (for reasons discussed earlier) that the blockchain have some novel features. Creating a blockchain with a novel feature also ensures that innovation continues to happen in the blockchain space at a much more fundamental level. It is clear that creating a new cryptocurrency is far more challenging than creating a new crypto-token.

Another important thing to note is that it is likely to be easier to spot a fraudulent offering of a new blockchain than to spot a fraudulent offering of a crypto-token. That is because the blockchain community is far more interested in developing a powerful blockchain; a powerful blockchain can support many powerful distributed applications but not the other way round. Hence, any blockchain that promises to have novel features is far more likely to be closely scrutinized than an ICO for a distributed application.

Hence, the SEC may find it beneficial to add a 'criterion of implementation' in addition to the criteria of the Howey test to identify whether a crypto-entity is a security. If the entity will be implemented as a cryptocurrency native to a *public* blockchain, then it is unlikely to be a security; otherwise, it is likely to be classified as a security. Such an interpretation will be in-line with existing legal precedents that a currency is not a security.

This raises the question of whether a native cryptocurrency on a private blockchain is a security? Again the implementation criterion is useful here since making changes to a privately controlled blockchain does not need a consensus among its participants whereas altering a public blockchain needs public consensus. Hence, the value of the cryptocurrency can be greatly altered by the entity that controls the private blockchain. Because such a scenario has potential for unfair market-practices, the SEC may find it useful to classify privately implemented cryptocurrencies as securities so that the investors are safeguarded from those unfair practices.

From the above discussion, it appears that the Kin Foundation's potential legal conflict with the SEC may be a rather landmark one for the future of cryptocurrencies and crypto-tokens. While this dissertation is not in a position to adjudicate on legal matters, it has certainly attempted to provide additional clarity on the underlying legal issue. This clarity may be useful for blockchain community, policy makers, and for law enforcement officials when they are making their decisions.

The first point to note from the above discussion is that it is clear that the core of the confusion emanates from functional-similarity of crypto-tokens to cryptocurrencies. Secondly, the functional-similarity of cryptocurrencies and crypto-tokens do not necessarily imply their technical-similarity as there is a significant difference in how and at what level in the technical protocol they are created. Lastly, it shows that Howey test may not be a sufficient test to assess whether a token offering can be classified as a security – although it may be a necessary test because of its relative simplicity and strong legal precedence. Adding a criterion of implementation may help clarify the difference between crypto-token and cryptocurrency, and thus make it easier to classify a crypto-token as a security or not while ensuring cryptocurrencies (native to public blockchains) are not classified as securities.

Another important trend of security tokens that may have implications for monetary policy is the 'tokenization' of traditional assets. Traditional assets like real estate assets are being converted to security tokens using blockchain. For example, the equity of a real estate entity is converted into a security token and then traded on the blockchain. Home owners may use this as a way to get more cash (or liquidity) from the equity they own on their homes. Since it is rather easy to fractionalize a security token on the blockchain, home equity can be broken up into multiple pieces and owned by different investors. This may result in an increasing trend of liquidating home equity i.e. increase the liquidation of traditional assets just by tokenzing them, and trading them on the blockchain. If monetary policy makers take liquidity into consideration in their efforts of managing inflation in the economy, then blockchain-led tokenization of traditional assets may be an important factor to consider.

4.3 Summary

This chapter showed that blockchain-created value can be interpreted in two ways: a) value created *by* blockchain, and b) value created *on* blockchain. Value created by blockchain comes from its usefulness as a market-maker, and as a secure and immutable record keeper. These categories of usefulness of blockchain in turn produce some generic policy implications like the potential increase in the number of illicit transactions and privacy violations (which will be addressed in the next chapter).

Value created on blockchain was shown to have rather complex policy implications, which do not seem to have easy or straightforward policy solutions. The main policy implications of this category can be divided into monetary policy implications and regulatory implications. The monetary policy implications come from cryptocurrencies, and an increasing trend of blockchain based tokenization of traditional assets (like real estate for example). If cryptocurrencies are widely adopted, it may affect the efficacy of traditional monetary policy – it is difficult to say, at this stage, in what way will it affect it because the transmission mechanisms are hard to trace out. The increasing tokenization of traditional assets is likely to contribute to increased liquidity in the economy just due to the act of tokenization. However, this seems to be far fetched issue as of now and may not need further investigation.

The regulatory implications of security tokens and cryptocurrencies are perhaps the hardest issues to address and also the ones that policymakers are rightly focused on addressing at the time of this writing. It is clear that the Howey test used by the SEC to determine if a token offering is a security may not be sufficient but it appears to be necessary test. The exploration in this chapter recommends that regulators must add a ‘criterion on implementation’ to the Howey test to clarify the differences between a crypto-token and cryptocurrency. Hence, it may help blockchain community, policymakers, and law enforcement officials in making more informed decisions on whether a crypto-entity is a security, and in what cases a cryptocurrency can be classified as a security.

5. Selected Policy Considerations

Essentially, a collective response to
a collective fear is: public policy.

In the previous chapters, this dissertation explored what is blockchain, how does it work, and what are the major changes that blockchain is bringing in. In closely examining the major changes that are occurring due to blockchain, many *potential* policy issues associated with those changes were flagged. This chapter proposes to dive deeper into those policy issues to explore what kinds of policy responses may be needed, if at all, to address those issues.

Let us start with a brief summary of the flagged policy issues. The first set of potential policy issues appeared to be associated with blockchain-enabled transactions of illicit good or services like drugs, illegal firearms, and child pornography. The second set of potential issues appeared to be associated with illicit methods of transaction like manipulating prices of goods or services for trade-based money laundering in private barter networks. Also, anti-competitive behavior of transacting parties in closed networks that exclude other parties fall in that category too. The next set of issues appeared to be associated with the purpose of blockchain-enabled transactions of crypto-assets due to their quasi-anonymity. For example, terrorism funding through transactions in cryptocurrencies, tax evasion and money laundering through undeclared investments in crypto-assets, and investor fraud through fraudulent ICOs.

Most of the policy issues posed by blockchain, as described above, are because blockchain can be used to:

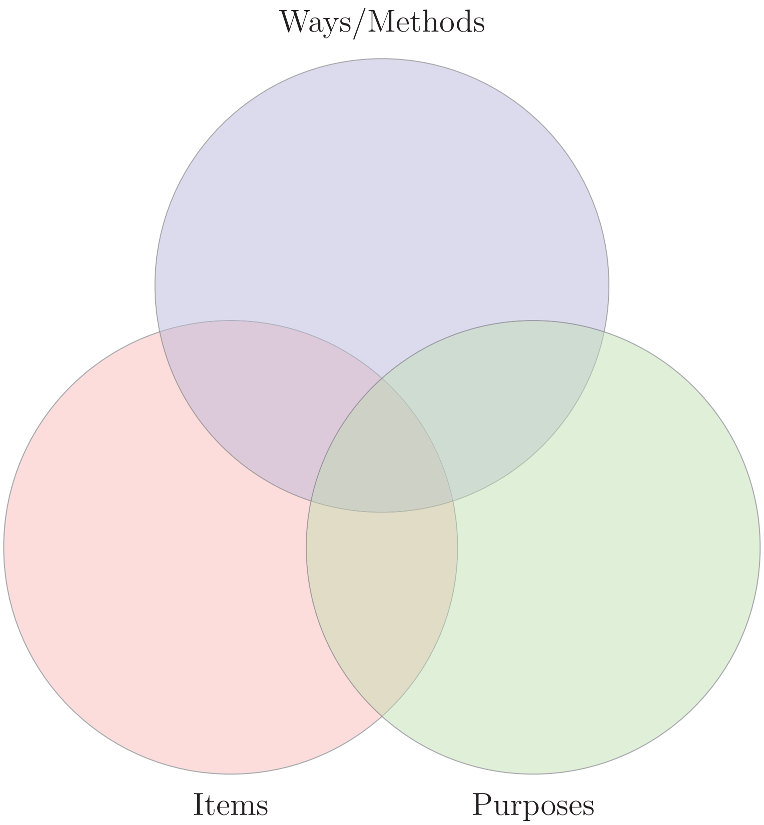

Figure 5.1: An Abstract Framework for Categorizing Policy Issues of Blockchain

- Transact illicit items
- Transact in illicit ways
- Transact for illicit purposes

Of course, blockchain can also be used to transact legitimate items, in legitimate ways, for legitimate purposes. There are some policy issues associated with those legitimate use cases as well. For example, using blockchain for sharing healthcare information between healthcare providers is an excellent example where even if everything is done legitimately, it may still pose issues associated with privacy of patient health information. This chapter focuses on the former set of policy issues since the latter ones are dependent on the specific application of blockchain.

5.1 Transaction of Illicit Items

As described in the chapter on intermediaries, the central policy puzzle associated with transaction of illicit items using blockchain is the accurate identification of who traded what items. This section attempts to explore the following questions about what can be done about transactions of illicit items using blockchain: a) Can law enforcement officials link transactions of illicit items to the identities of transacting parties? b) Can official intercept such transactions? and c) Can they disrupt such transactions?

5.1.1 Linking Transactions with Identities

There appears to be high linkability of transactions carried out using bitcoin-blockchain. Fleder, Kester, and Pillai (2015) show that by using publicly available information about transacting parties and the transaction ledger from the bitcoin-blockchain one can link transactions to online names of users. Of course, 'online names of users' here refers to the name that was provided by the user while creating an account online. This implies that if someone with a legitimate account at a legitimate financial institution wants to engage in illicit transactions using bitcoin, their identity can be linked to those transactions. As long as an user name can be traced back to a real identity and transacted bitcoins are eventually converted into a fiat currency, there is a high likelihood of linking transactions with the identities of transacting parties.

Because of such high linkability of bitcoin-based transactions, many illicit activities tend to be carried out using cryptocurrencies like Monero, Zcash, and others that provide higher privacy compared to bitcoin. However, it appears that even such transactions can be traced. Kumar et al. (2017) through their forensic analysis show that by exploiting certain heuristics of Monero's blockchain, they were able to trace 87% of the transactions. It should be mentioned, however, that of the 87% of transactions 65% of transactions were by default traceable. So their forensic analysis enabled traceability of an additional 22% of transactions that were not by default traceable. This means that of the 35% of transactions that did not want to be traced, 22% were traced using their forensic analysis. This implies that there are transactions on some blockchains that are hard to trace. However, this will be an ongoing effort to trace transactions on such blockchains. The US Department of Homeland Security is apparently seeking to address this issue as shown by allocating seed-funding start-ups that are interested in developing innovative forensic analytics solutions to the problem of tracing transactions involving Monero and Zcash.[1] This shows that it is an important national security concern, and there are research efforts underway to address this issue. It also shows that policy makers are trying to address this issue in a way that is boosting innovation, rather than clamping it down.

The main hurdle to be crossed in forensic analysis of transactions on blockchain

[1] See `https://www.sbir.gov/sbirsearch/detail/1547625` Accessed May-2019

is that of finding the needle in the haystack. The sheer scale of transactions to be analyzed to find the problematic ones may make the task rather difficult. Hence, researchers have tried applying machine learning algorithms that automate the process of finding the needle from a big haystack. Sun Yin et al. (2019) show that supervised machine learning techniques can be used for de-anonymizing transactions on the bitcoin-blockchain. They were able to achieve a cross-validation accuracy of roughly 80%, which is quite good. They also claim that their algorithms can be used to estimate the total amount of cybercriminal activity that uses bitcoin-blockchain, since the algorithm has also been trained to identify the network of transactions involved in cybercriminal activities.

It remains to be seen if the same machine learning model can be used to identify criminal activities on other blockchains as well apart from bitcoin-blockchain. There is some progress being made on this front as well. Graphsense is an open-source platform being developed by the Austrian Institute of Technology that can be used to analyze transactions across multiple ledgers.[2] Such a cross-ledger analytics platform can be quite useful in identifying illicit transactions that span multiple blockchains. But as of this writing it does not seem quite easy to identify such transactions that involve multiple blockchain transactions.

However, wherever transactions across multiple blockchains are involved there has to be a crypto-exchange to enable the cross-currency transaction. Since registered crypto-exchanges are bound by KYC regulations as laid down by the SEC, it may deter criminal entities from using registered crypto-exchanges for obvious reasons. There are, however, some decentralized exchanges that do not require their customers to comply with Know Your Customer (KYC) and Anti-Money Laundering (AML) regulations. For example, EthFinex Trustless appears to be one such exchange that does not talk about any AML or KYC compliance on its website.[3]

[2] See `https://graphsense.info/` Accessed May-2019

[3] See `https://trustless.ethfinex.com/` Accessed May-2019

That said, EthFinex has a section that claims to cooperate with law enforcement agencies if they follow proper procedures as laid down in the jurisdiction the agency operates in.[4] Most other decentralized exchanges like IDEX that offered services without KYC/AML regulations are reportedly now trying to become fully compliant with those regulations.[5]

That said, the mechanism of "atomic swap" or peer-to-peer swap that is being gradually adopted is an interesting case that seems to have the potential to remove the need for crypto-exchanges. Atomic swaps are like smart contracts that two users wanting to directly swap their cryptocurrencies with each other can use on their separate blockchains to execute that swap. Although it sounds like no intermediaries are needed, there needs to be an entity that has access to both blockchains. Companies like Liquality are providing such services.[6]. Interestingly their policy is that "[d]own the line we will enable optional KYC integrations for institutions who need to meet specific compliance measures, while assuring that an alternative remains in place for everyone."[7]

All of this suggests that policy makers and law enforcement officials are aware of blockchain based transactions being used for illicit activities, and they are not lagging far behind in regulating its use for such activities. Also, the fear of regulatory enforcement by bodies like the SEC is forcing crypto-exchanges – decentralized or not – to comply with regulations (though atomic-swap services companies, as of this writing, seem to be assuring that they will continue to support the option of complete decentralization of currency exchanges). Moreover, policy makers are trying to address the issues of de-anonymizing transactions in an innovation-friendly manner by providing seed-funding for start-ups willing to address those issues.

[4] See `https://www.ethfinex.com/legal/law_enforcement_requests_policy` Accessed May-2019

[5] See `https://news.bitcoin.com/decentralized-exchange-idex-to-introduce-full-kyc/` Accessed May-2019

[6] See `liquality.io` Accessed May-2019

[7] See `https://liquality.io/faq/\#registerWithLiquality` Accessed May-2019

5.1.2 Intercepting Illicit Transactions

Broadly speaking, there are two points of contact for intercepting illicit transactions. The first is to intercept the payment side of the transaction and the second is to intercept the delivery side of the transaction. To intercept the payment side of the transaction, authorities need to de-anonymize the transaction as discussed above. To intercept the delivery side of the transaction, authorities need to know the nature of the goods or services being delivered.

One of the issues highlighted in the chapter on intermediaries was that the USPS is prohibited from opening a package that is being routed through their services. Private delivery services like FedEx and others do have the legal basis to open a package to inspect it if they find it suspicious. Because of such legal restrictions, it is possible that illicit items are being delivered through USPS (but also through other private delivery services).[8] To combat this issue policy makers may find it useful if authorities in the USPS had non-invasive inspection capabilities – those that can inspect the nature of the items being shipped but without opening the package.

While variants of such technology are employed at airports, it is not completely clear (perhaps for reasons of confidentiality) whether USPS is using such technologies. However, according to the Office of Inspector General's report of 2018, the data-analytics capabilities of USPS needs to be enhanced. Specifically, the data model that USPS has built over time apparently lacks information about whether a package flagged for illicit items resulted in seizure of those items by law enforcement

[8] See https://www.dea.gov/sites/default/files/2018-11/DIR-032-18%202018%20NDTA%20final%20low%20resolution.pdf Accessed May-2019

authorities.[9] This lack of information on successful seizure of those flagged packages renders the model weak.

Since cryptocurrencies are likely contributing to the increase of transactions involving illicit items, information sharing about successful seizure between law enforcement officials and USPS officials becomes even more important, if illicit transactions are to be effectively intercepted. This suggests that even though a lot of effort is being expended on identifying illicit transactions, its benefits are not being fully derived. This is an important policy gap highlighted by increasing illicit transactions that rely on blockchain and cryptocurrencies.

5.1.3 Disrupting Illicit Markets

The next question to be explored in this domain is whether illicit markets enabled by blockchain can be disrupted. To disrupt an online market for illicit items policy makers can, among other things, employ a combination of the following options:

- Seize substantial quantities of illicit items
- Nab the dealers and sellers of those illicit items
- Shutdown the online marketplaces where these transactions occur
- Deter the buyers from buying those illicit items
- Cut-down the payment channels to sellers
- Cut-down the delivery channels to buyers

Evidence suggests that law enforcement agencies are trying to do every single one of the above set of tasks. The DEA is trying to seize substantial quantities of il-

[9] See pg.14 Finding #4 in the Office of Inspector General's report `https://www.uspsoig.gov/sites/default/files/document-library-files/2018/SAT-AR-18-002.pdf` Accessed May-2019

licit items.[10] They are also constantly trying to bring drug dealers to justice. Many marketplaces have been shut down. For example, Silk Road (one of the biggest drug market online) was reportedly shutdown in 2013 after FBI arrested its creator Ross Ulbricht,[11] and another big drug market called Dream Market mysteriously shutdown its operations in April 2019.[12] Since crypto-exchanges require their customers to comply with AML and KYC regulations, even buyers are likely to be deterred from engaging in illicit transactions with the cryptocurrencies they have bought from exchanges because most of the information is increasingly traceable if law enforcement officials want to trace the transaction. Law enforcers are also trying to make it harder to deliver illicit goods using existing delivery mechanisms, as discussed in the previous section.

The only task that blockchain has made it harder for law enforcement officials is to cut down payment channels for payments being made using cryptocurrencies. Payment channels and dark-web marketplaces for illicit items have only become more resilient with the advent of blockchain and cryptocurrencies. As long as the supply of illicit items continues, there is always a place to buy those items on the dark web. Décary-Hétu and Giommoni (2017) show that police crackdowns on illicit drug marketplaces affect those markets only for a short duration of time. The crackdowns do not necessarily affect those markets in the long term because new marketplace pop-up every month and it is easy for participants to move from one marketplace to another. They showed that due to policy crackdowns "[b]oth the supply of and the consumption of drugs were impacted, though drug prices appear to have remained unchanged." They conclude that police crackdowns were ineffective at lowering

[10] See https://www.dea.gov/sites/default/files/2018-11/DIR-032-18%202018%20NDTA%20final%20low%20resolution.pdf Accessed May-2019

[11] See https://www.forbes.com/sites/andygreenberg/2013/10/02/end-of-the-silk-road-fbi-busts-the-webs-biggest-anonymous-drug-black-market/#6aad94e95b4f Accessed May-2019

[12] See https://www.vice.com/en_us/article/wjmw3w/the-worlds-biggest-dark-net-market-has-shut-downwhats-next-1 Accessed May-2019

the volume of sales on online drug markets. Similar conclusions were arrived at, independently, by Ladegaard (2017) and Cunliffe et al. (2017).

This implies that online markets for illicit items are more resilient to law enforcement actions, and that blockchain and cryptocurrencies are only making them more resilient to such actions. Given such resilience it may be harder to disrupt such illicit markets. Hence, it only means that identifying illicit transactions, linking those transactions to perpetrators, and intercepting those transactions are even more important, if policy makers want to keep a check on illicit markets enabled by blockchain and cryptocurrencies.

5.2 Transacting in Illicit Ways

As discussed earlier, blockchain has the potential to change the way businesses trade with each other and their own subsidiaries (within and across borders). When businesses change the way they interact with each other especially when linked with a blockchain, it can, among other things, increase the likelihood of anti-competitive behavior. When businesses change the way they interact with their subsidiaries while still being linked through a blockchain, it raises concerns about transfer pricing. This section attempts to delve deeper into these potential issues to examine closely if and how blockchain is fueling those policy concerns, and what could be done to address those concerns.

5.2.1 Anti-Competitive Behavior

As seen in the case of Walmart bringing in some of it suppliers and vendors under a common blockchain for a pilot project to simplify the their supply chain process, blockchain could lead to a general trend of groups of businesses coming together under one blockchain. Blockchain has the potential to enable what can be called

Islands of Commerce on the Internet. In fact, one such island seems to be in the making.

In June 2019, Facebook announced that it was launching a cryptocurrency (and an associated blockchain) called Libra.[13] Libra is proposed to be used as a global cryptocurrency that "will support peer-to-peer payments and a few other ways to pay such as QR codes [(Quick Response codes)] which small merchants can use to accept payments in Libra. Over time there will be many others including in-store payments, integrations into Point-of-Sale systems, and more."[14]

While such cooperation between businesses is a welcome development because of its potential to promote efficiency (of payments in this case), it could potentially give such a group a substantial market-advantage over market-players who are not part of the group. Any exclusionary conduct by such a group could lead to anti-competitive practices.

Thankfully, Section 5 of the Federal Trade Commission Act in the US has provisions that ban "unfair methods of competition" and "unfair or deceptive acts or practices." These laws are written in a technology-agnostic way. They rightly attack the problem of anti-competitive behavior rather than the means or the technologies that enable such a behavior.

Inherently, blockchain has nothing in it that makes it anti-competitive in nature. If anything, blockchain is a technology that is enabling competition in areas where the competition is quite limited. For example, governments have had the monopoly over issuing new units of their national currency for several decades now. With the advent of blockchain, groups of companies – as seen in the case of Libra – are coming together to issue a currency that can be accepted across borders unlike fiat currencies.

Thus, as long as the entities using blockchain to form a digital island of com-

[13] See https://libra.org Accessed June-2019

[14] See the FAQ section on https://calibra.com/ Accessed June-2019

merce on the Internet are governed by anti-competitive laws, the risk of blockchain exacerbating anti-competitiveness remains trivial. That said, it may be worth discussing a case related to anti-competitiveness that is currently (as of this writing) pending before the District Court in Southern District of Florida. This case could provide some insights into what may anti-competitive behavior mean in the context of blockchain.

In the case of *United American Corp. v. Bitmain Inc. et al.*, the plaintiff United American Corp. (UAC) alleges that the defendants (Bitmain) colluded with other market-players in manipulating the currency market for Bitcoin Cash. To understand that allegation, one must begin with understanding what is a fork on a blockchain.

A fork on a blockchain is essentially two (or more) different chains of blocks (of transactions) that emerge from a single chain of blocks within a blockchain. A fork that is short-lived or temporary is seen as a soft fork,[15] and one that is more permanent is known as a hard fork. There are multiple reasons why more than one chain of blocks could emerge from a single chain of confirmed blocks on a blockchain.

In the UAC v. Bitmain case, the hard fork on the Bitcoin Cash blockchain emerged in November 2018 because there was disagreement among the miners of the blockchain regarding a technical update to the blockchain protocol for various reasons.[16] One prong of the fork was called Bitcoin ABC that the group of miners led by Bitmain *et al.* was in favor of and the other prong was called Bitcoin SV (Satoshi's Vision) that mining group led by UAC was in favor of.

When a hard fork emerges on a blockchain, the miners on the blockchain have to choose one of the various forked chain of blocks to continue their mining activities on. According to UAC, many of the entities in the mining communities were influenced

[15] It can happen due to what computer scientists call a 'race-condition' on the blockchain

[16] See `https://www.investopedia.com/news/all-about-bitcoin-cash-hard-fork/` for more details. Accessed June-2019

by Bitmain Inc. *et al.* to choose the chain of blocks that Bitmain *et al.* were in favor of (i.e. Bitcoin ABC). In the end, Bitcoin ABC gained the most number of miners and succeeded as Bitcoin Cash – not to be confused with the original Bitcoin. UAC alleges that Bitmain *et al.* used anti-competitive practices in taking over the Bitcoin Cash network. Since the case is pending before the court, it may best be left to courts to adjudicate on this matter. However, teasing out a general trend that may have implications for policy makers may be a useful endeavor.

Hard forks due to technical updates to blockchain protocols are likely to be quite common, since the technology is evolving continuously. If the stakes are high for the miners involved with hard-forked blockchains, law enforcement officials will likely see an increasing number of cases alleging anti-competitive behavior similar to the UAC v. Bitmain case. It is not difficult to see that such cases may not strictly fall under "exclusionary conduct," since UAC was not excluded by Bitmain *et al.* from participating in the Bitcoin ABC fork. It may be more accurate to classify such cases under what can be termed as "isolationary conduct" – whether the conduct was illegal or not remains to be seen. Since such cases are likely to emerge regularly, it may be helpful for the courts to decide on such matters if the FTC adds some new clauses to the existing acts defining what is "isolationary conduct", and when is it illegal if entities that previously belonged to a group are left isolated after certain updates to the rules of the market they are participating in.

That said, it should be noted that there is no proper boundary defining the meaning of exclusionary conduct. As the former Director of the Bureau of Competition at the FTC, Susan Creighton, had put it in 2005: "it is a conceptual problem that has not yet been fully solved – at least to the extent that there is no consensus on a universal test for determining exclusionary conduct."[17] A panel discussion held in

[17] See pg. 2 of "Ranking Exclusionary Conduct" https://www.ftc.gov/sites/default/files/documents/public_statements/ranking-exclusionary-conduct/051115conduct.pdf Accessed June-2019

October 2018 as part of the third Hearing on Competition and Consumer Protection shows that there is still no clear definition of exclusionary conduct.[18] Moreover, the discussion seems to suggest that the definition has gotten more complicated with multi-sided platforms. Not a lot appears to have changed (for the better) since 2005.

However, it may be beneficial to think of exclusionary and isolationary conducts in terms of gaining and losing access to a market. Exclusionary conduct can suggest that an entity could not gain access to a market due to illicit methods employed by its competitors. On the other hand, isolationary conduct can suggest that an entity loses access to a market due to illicit methods (for example, collusion with other entities) employed by its competitors. Obviously, both of those conducts can affect healthy competition in any market.

Anti-competitive cases involving private “permissioned” blockchains (like the one being piloted by Walmart, for example) can involve either exclusionary or isolationary conducts whereas cases involving public “permissionless” blockchains (for example, the Bitcoin Cash blockchain) can likely only involve isolationary conduct. The main point here is that exclusionary conduct does not make a lot of sense in the context of open source, public, and permissionless blockchains (or markets). Hence, it may provide more clarity to law enforcers if FTC closes the legal gap by including isolationary conduct to the existing acts that combat anti-competitive behavior.

5.2.2 Transfer Mispricing

Generally speaking, transfer pricing comes into play when a larger multi-entity firm conducts business with its individual entities in a way that treats and measures them as separate entities. When entities transact goods and services with each other under the umbrella of a larger multi-entity framework, the prices used for those goods and

[18] Listen to the panel discussion at `https://www.ftc.gov/news-events/audio-video/audio/understanding-exclusionary-conduct-cases-involving-multi-sided` Accessed June-2019

services are called transfer prices or transfer costs. Blockchain is a technology that is appearing to be effective as a tool for managing a unified ledger for transactions between individual entities of a multi-entity organization or between the entities of a consortium. This shows that both blockchain and transfer pricing can interact with each other under a multiple entities framework. Since transfer pricing is an important policy issue, it may be useful to explore this topic in greater detail.

One of the main concern policy makers have with respect to transfer pricing is the evasion of taxes through transfer mispricing (or manipulation of transfer prices). If transacting entities are related to each other, there is an incentive for them to transact goods and services with each other but also to price those goods and services in a way that reduces their individual tax liabilities. When two entities are not related to each other, the prices they use to transact with each other can, by and large, be considered as fair market prices. However, when the entities are related to each other the fairness of prices may be subject to suspicion because of their incentives to manipulate those prices.

It is quite plausible that with the advent of blockchain based islands of commerce, commercial entities that may not be related to each other can still cooperate with each other in a way that may provide the trust, the avenues, and the incentives they would need to manipulate prices. This is because members of a digital island of commerce may risk isolation if they act counter to the interests of the island (or consortium). Only trustworthy entities are likely to be members of the consortium held together by a blockchain.

Secondly, since the market-maker is the market in the case of blockchain, buyers and sellers can easily fulfill each others' commercial needs for good and services. This shows that blockchain provides the avenues for buyers and sellers of a consortium to transact with each other.

Thirdly, if the buyers and sellers in the consortium can reduce their tax liabilities by cooperatively pricing their goods and services in away that is beneficial to each other (and to the consortium, if needed), it can be enough of an incentive to misprice their goods and services without compromising on the value derived from them.

Lastly, because of the nature of blockchain, the more transactions occur between the members (and customers) of the consortium the more units of cryptocurrency are generated (or come into circulation) as mining reward for the consortium. This may provide more incentives for them to transact with each other and with each others' customer base. That said, it may also contribute to inflation of their cryptocurrency – but that is likely something that the consortium has to deal with; not policy makers necessarily. It may be useful to point out that while the trust, avenues, and cooperation-incentives mentioned above may apply to a consortium based on a barter network as well, the mining reward incentive is only applicable to a consortium based on blockchain. This indicates that members of a consortium that uses blockchain may have incentives to transact with each other even if they are not related to each other as individual entities of a multi-entity firm.

Thus, the fundamental implication of blockchain for transfer pricing is that it expands the scope of transfer pricing from related entities under a multi-entity firm to cooperating entities linked together with a blockchain. It implies that tax agencies may have to cover a larger ground to ensure blockchain linked commercial entities are not evading taxes through transfer mispricing.

The way IRS in the US has sought to address the issue of tax mispricing in barter networks is by mandating the barter exchanges to submit Form 1099B titled "Proceeds From Broker and Barter Exchange Transactions" for every participant on the barter network.[19] Every participant's Form 1099B totals the sales through barter for that participant. There is no reason not to use the same method or to extend the scope of this method to blockchain linked entities as well. Mandating the

[19] See `https://www.irs.gov/pub/irs-pdf/f1099b.pdf` Accessed June-2019

entity that manages the blockchain that links various entities into a consortium to submit a Form 1099B (or a similar form) may be enough in terms of addressing this policy gap. That said, one must note that addressing the policy gap and enforcing compliance to those policies are two significantly different things.

Another important policy aspect with respect transfer pricing in the context of blockchain linked commercial entities that needs to be examined is: whether the methods used by regulatory agencies to calculate market prices of goods and services are robust enough to be used in the case of blockchain linked commercial entities as well?

The Organization for Economic Cooperation and Development (OECD) has laid down five different methods for calculating transfer prices.[20] Most government agencies around the world are known to employ these methods – as appropriate – for calculating transfer prices. It may be useful to begin with briefly summarizing those methods before exploring if any changes may be needed for calculating transfer prices in the context of islands of commerce enabled by blockchain.

OECD recommends five different methods for calculating transfer prices. They are divided into two categories: a) traditional transaction methods, b) transactional profit methods. Put simply, the first category relies more on prices to calculate transfer prices and the second category relies more on profits to calculate transfer prices.

Transaction Methods

1. *Comparable Uncontrolled Price (CUP) Method*: The CUP method involves comparing the terms and conditions of a "controlled transaction" with that of a "uncontrolled third party transaction." – these terminologies are applicable for all the

[20] See `https://www.oecd.org/ctp/transfer-pricing/45765701.pdf` Accessed June-2019

following methods. A controlled transaction is a transaction between two individual entities under a multi-entity firm. On the other hand, an uncontrolled third party transaction is a transaction between two unrelated entities. CUP method uses the price of a similar third party transaction as a reference for calculating the price for a controlled transaction.

2. *Resale Price Method*: This method involves calculating the 'resale price', calculating the 'gross margin' on resale price, and then taking the difference between the two to arrive at the transfer price. The resale price is the price at which an "associated entity" – an entity that does not belong to the multi-entity firm – sells the goods/services to a third party. The gross resale price margin is the margin maintained by the associated entity. Thus, $TransferPrice = ResalePrice - ResalePriceMargin$.

3. *Cost Plus Method*: This method involves calculating the 'cost of sales', calculating the appropriate market-based 'mark-up', and adding them together to arrive at the transfer price. The cost of sales is the total costs incurred by a supplier involved in the controlled transaction with a related entity. The mark-up is calculated by comparing the gross profits and cost of sales associated with an uncontrolled third party transaction on the market. Thus, $TransferPrice = Cost + Mark\_Up$.

Profit Methods

4. *Transactional Net Margin Method*: This method involves calculating the net profit of controlled transactions, calculating the net profits of uncontrolled transactions, and comparing the terms and conditions to arrive at appropriate transfer prices. The method specifies that the comparable uncontrolled transaction can be

"broadly similar" if an exact match is not found.

5. *Profit Split Method*: Sometimes individual entities of a multi-entity firm transact with other entities in a rather interrelated way. For example, they may agree to set up a joint venture and share profits from that venture. The profit split method calculating the profits shared between controlled transactions, calculating the profits shared between uncontrolled third party transactions, and comparing their terms and conditions to arrive at an appropriate transfer price.

From the above brief summaries, it appears that the methods used for calculating transfer prices rely heavily on comparing controlled and uncontrolled transactions, and deriving an appropriate transfer price. While these methods have worked reasonably well for pricing the value created by traditional companies, these methods have been criticized for not being quite effective at pricing the value created by digital companies. In fact, that is one of the main factors driving OECD to develop the Base Erosion and Profit Shifting (BEPS) guidelines.[21]

Among the challenges created by digital companies, "network effects" is a prime contributor to the complications of transfer pricing. Network effect is broadly described as a phenomenon where goods or services offered by a commercial entity gains additional value as more people use them. This phenomenon makes calculation of transfer prices tricky because profits are not comparable between transactions that involve network effects and those that do not. Even transactions involving two different networks are quite hard to compare because their network effects are not easy to match.

Since blockchain is a technology that, by definition, involves creation of a network of entities that can transact with each other, network effects are likely to be native to blockchain enabled digital islands of commerce. Thus, any calculation of transfer

[21] See https://www.oecd.org/tax/beps/public-consultation-document-addressing-the-tax-challenges-of-the-digitalisation-of-the-economy.pdf Accessed June-2019

prices involving profits will somehow have to incorporate the impact of network effects. Incorporating network effects into transfer pricing is an active area of policy research with no clear answers yet. OECD is seeking to examine these issues closely by setting up a Task Force on the Digital Economy. OECD has also issued a public consultation document to collect responses and ideas from the public on this matter.

It appears that due to the difficulty of calculating transfer prices accurately, especially of a blockchain based consortium, it is plausible that at least some portion of taxes may not be fully recovered by tax agencies. Furthermore, since the value of goods and services is subject to change due to the particular network effects at play within a network of commercial entities, it can be difficult to state clearly (without proper investigation) whether the invoices are being properly valued or if they are over- or under-valued. This artifact could be exploited by malicious entities for illicit purposes like money laundering.

5.3 Transacting for Illicit Purposes

The above section shows that blockchain can enable entities (if they wish so) to transact in ways that are illicit, and those transactions can be quite difficult to be classified as illicit without time-consuming investigations by law enforcement agencies. This lack of clarity about the illicitness of transactions could be exploited by malicious entities to engage in transactions that serve illicit purposes. Money laundering and terrorism financing are two such purposes, among others, that have raised a lot of concern among policy makers. This section explores the role blockchain based transactions can play in enabling money laundering and terror financing.

5.3.1 Money Laundering

Money laundering refers to the process that conceals the origins of illegally obtained money by routing the money through a complex maze of transfers and commercial transactions. Passing money through such a complex maze of transactions obscures its origins, and makes the money seem like it was obtained legally. Broadly speaking, there are three steps that are commonly used by money launderers: a) placement, b) layering, and c) integration.

Placement refers to the point where the illegally obtained money is furtively introduced into a legitimate financial system – usually in countries that are known to be tax havens because not a lot of documentation for money is mandated in those countries. Layering refers to the process of passing the money through a complex maze of transactions that creates a lot of confusion about the trail and ownership of that money. The money is passed through multiple accounts, exchanged into multiple currencies, and used in multiple legitimate transactions before being integrated into the accounts of the laundering entity. Successful integration makes the money seem like it was obtained by the (laundering) entity through legitimate purposes.

Some of the questions with respect to blockchain and money laundering that may produce important policy relevant insights are: What does the interaction between blockchain and money laundering look like? Does blockchain make it easier to launder money? If so, how?

To explore the above question, one can begin with looking at where may the utility of blockchain lie in the three broad steps of money laundering discussed above. Based on how blockchain enables transactions between two or more entities, it appears that blockchain may have the most utility in the 'layering' step of money laundering. Using blockchain based transactions (involving multiple fiat and crypto-assets), an entity can create quite a complex maze of transactions that can be hard to trace, if not impossible. Note that the difficulty of tracing the origins of money does not necessarily come from lack of transparency of transactions – in fact, there is high transparency and traceability of transactions on blockchain. The difficulty comes from the law enforcers' inability to access information that can help them link transactions to the identities of transacting parties. With blockchain based transactions, transparency of transactions is high but also privacy of transacting entities is high due to quasi-anonymity of transacting parties.

As long as the KYC regulations are strongly complied with by risk-absorbing intermediaries operating in the blockchain space, law enforcement officials may be able to link the identities of transacting parties with transactions with proper legal authorizations. Linking identities with transactions helps law enforcement officials in navigating the complex maze deliberately created by the laundering entity to obscure the trail and ownership of illegally obtained money.

That said, the blockchain space does offer some features that can be potentially exploited by laundering entities to make the process of navigating the maze quite difficult for law enforcement officials. For example, as discussed earlier, there are some cryptocurrencies like Monero and ZCash that have high privacy protecting features. Details of transactions involving these cryptocurrencies may be harder to piece together than other cryptocurrencies like Bitcoin.

In addition, there are "tumbling services" or "mixing services" on offer that randomly mix potentially identifiable cryptocurrency funds with other funds in such a way that it obscures the trail back to original source of those funds. This random mixing, the service providers claim, breaks the chain of transactions that can link the crypto-coins they sent in with the crypto-coins they received after the random mixing.[22] In return, the service providers tend to charge a certain percentage of coins

[22] See `https://smartmix.io/faq` for example. Accessed Jun-2019

per mixing service and a certain fixed fee per receiving address the user provides to receive the coins after they are randomly mixed.

There are many such tumbling services being offered at the moment but not many of them talk about complying with KYC regulations yet (at the time of this writing). However, regulators have been making important moves to ensure these tumbling services properly comply with existing regulations. The recent guidelines issued by the Financial Crimes Enforcement Network (FinCEN) suggest that tumbling or mixing services are nothing but newer forms of "money transmission services" that engage in transmitting cryptocurrencies. Money transmission service is defined as "the acceptance of currency, funds, or other value that substitutes for currency from one person and the transmission of currency, funds, or other value that substitutes for currency to another location or person by any means" (FinCEN, 2019). As such, tumbling or mixing service providers are subject to Bank Secrecy Act (BSA) and AML regulations. This is yet another instance that shows existing policies and regulations are robust enough to handle changes brought in by new technologies in the financial domain.

This suggests that as long as the tumbling service providers comply with BSA and AML regulations as applicable to money transmission services, law enforcement officials should be able to link the identities of entities availing tumbling services to obscure the source of their funds – regardless of the chain of transactions being broken by random mixing. This is because the regulations, if complied with fully, can ensure that the identity of the entity availing the tumbling services is recorded such that it can be linked with the sending and receiving addresses. This ensures that those who are availing those services to protect their privacy rather than to conceal their illegal activities will be able to enjoy enhanced privacy. Whereas those using the tumbling services to launder illegally obtained money may find it difficult to succeed at laundering that money.

Of course, the scale of the money laundering problem is growing larger quantitatively due to blockchain and cryptocurrencies; however, the nature of the problem has not changed qualitatively. Thankfully, the ability to scale-up or to automate existing solutions to those problems is also growing due to newer methods that can be developed using AI (see (Sun Yin et al., 2019), and (Chang and Svetinovic, 2018)). Thus, even though the scale of the problem is growing, the solutions also have the potential to scale up.

5.3.2 Financing Organized Crime and Terrorism

The above discussion focuses on bringing illegally obtained money into the financial system so that it can be used without raising concerns of law enforcement officials. There is also another kind of manipulation of the flow of money that channels legally obtained money out of the financial system to fund criminal and terrorist activities. This form of channeling legally obtained money for illicit activities is done through what is known as trade based money laundering (TBML).

TBML involves mipricing the invoices such that the invoices are over-valued but the actual money spent is less than the value quoted in the invoice (similar to transfer mispricing). The difference in value can then be channeled to fund illicit purposes. As discussed earlier, mispricing is quite plausible in the case of blockchain linked entities because calculating the prices accurately when network effects are present can be quite difficult.

Moreover, in the case of blockchain and cryptocurrencies creating an address to receive money is quite easy and does not need any KYC compliance. Entities can choose to create a new and unique address – tied to a specific cryptocurrency – for every transaction, if they wish to. Thus, receiving cryptocurrencies is quite easy.

It has been reported that some terrorist organizations are using this technique to safely raise funds to support their activities.[23]

Some academic studies have also explored the topic of how could terrorists exploit cryptocurrencies to sustain their operations. Irwin and Milad (2016) state that "[a]lthough it is difficult to find concrete evidence of largescale use of Bitcoins and other crypto-currencies by terrorist groups and their supporters, there is strong evidence to suggest that they have been linked to a number of terror attacks in Europe and Indonesia." Also, they used documents published by ISIS to create a set of models to show how a person interested in getting to Syria or Iraq to join the 'jihadists' could do so just using bitcoins.

That said, there is also quite a bit of skepticism about whether terrorists are actually using cryptocurrencies to fund their activities. A recent report published by RAND Corporation suggests that "[c]urrent cryptocurrencies are not well matched with the totality of features that would be needed and desirable to terrorist groups..." (Dion-Schwarz, Manheim, and Johnston, 2019). However, it also forecasts that "should a single cryptocurrency emerge that provides widespread adoption, better anonymity, improved security, and that is subject to lax or inconsistent regulation, then the potential utility of this cryptocurrency, as well as the potential for its use by terrorist organizations, would increase." Thus, it is not completely clear whether terrorist organizations are using cryptocurrencies to fund their operations.

Nevertheless, there is substantial evidence to show that TBML has been used to fund terrorist organizations (Zdanowicz, 2009). Arguably, if blockchain can enable TBML then it can, at least in theory, indirectly enable channeling of funds for terrorist and organized criminal activities. However, if blockchain is exploited for TBML then it is likely to be in a setting where multiple unrelated commercial entities come together to cooperate with each other to reap the benefits of network effects.

[23] See https://thehill.com/opinion/cybersecurity/377415-terrorists-have-been-using-bitcoin-for-four-years-so-whats-the-surprise Accessed July-2019

In such a case, the entities trying to engage in TBML to fund politically controversial causes like terrorism may risk isolation from the rest of the consortium. TBML for funding terrorism, arguably, may be easier to carry out in a smaller setting where only a couple of entities are involved. If the number of entities is larger, as is likely to be the case in a blockchain linked consortium, then it may be too risky politically for the entire consortium. Hence, it is plausible that forces within the blockchain linked consortium may be quite vigilant about illegal activities like TBML – especially so if it is being used to fund terrorism – to safeguard their own individual reputations. Therefore, it appears that if policy makers and law enforcement officials can combat money laundering the risk of terrorism funding enabled by blockchain is likely to be quite low.

5.4 Summary

This chapter began with recognizing that many of the generic policy issues raised by blockchain are related to blockchain being used for transacting illicit items, transacting in illicit ways, and transacting for illicit purposes. After analyzing those three cases in detail, we found that although blockchain has the potential to aggravate the associated policy issues, the existing policies themselves are quite robust to the changes brought in by blockchain. Moreover, policy makers are doing many things to close the loop-holes created by blockchain and cryptocurrencies in a manner that is innovation-friendly.

The major implication across all policy concerns raised by blockchain appears to be that law enforcement officials may likely have a quantitatively higher number of cases to pay attention to. Qualitatively, however, the nature of problems does not appear to have changed because of the advances made by blockchain. This implies that intelligent automation of solutions using AI may be quite helpful for law enforcement to be effective at fighting illicit activities enabled by blockchain and cryptocurrencies.

6. A Potential Application of Blockchain

Uniformity promotes scalability.

Diversity promotes security.

Cyber attacks tend to spread like wildfire because of one main reason: identical software. If an adversary successfully exploits a vulnerability on one system, then that adversary can likely succeed in exploiting that vulnerability on many other systems running on identical software. Attacks like WannaCry, NotPetya, and others have exploited identical nature of software to successfully compromise a large number of systems around the world. As of May 2017, the ransomware WannaCry had reportedly compromised at least 200,000 systems worldwide and demanded a certain payment in bitcoins to unlock the victim's computer.[1]

In the age of autonomous systems like drones, unmanned aerial vehicles (UAV), and self-driving cars where various software govern the movement of those machines, it is rather important to ensure that it remains difficult and expensive for an adversary to successfully exploit those systems in large numbers. If these systems are compromised *en masse*, the damages adversaries can cause using them can be quite substantial.

Of course, such events are likely to have a very small probability of materializing (like blackswan events). However, the damages such events can cause if they materialize can be catastrophic in nature. Since autonomous systems are still in their infancy (but are rapidly growing in numbers), it may be valuable to consider developing non-identical software with identical functionality (NISWIF) as a strategy to mitigate large scale ($n \gg 1$) exploitation of those systems.

[1] See https://www.abc.net.au/news/2017-05-15/ransomware-attack-to-hit-victims-in-australia-government-says/8526346 Accessed Apr-2019

6.1 Can NISWIF Mitigate Large Scale Exploitation?

An analogy may be useful in introducing how a NISWIF strategy can potentially mitigate large scale exploitation of autonomous systems. Consider human genes for example; if all humans had identical genes, then a disease that was capable of exploiting a genetic vulnerability could quickly turn into a pandemic. Thankfully, genetic diversity introduces a certain degree of uncertainty that ensures that not all humans are susceptible to an infection when exposed to a disease. Thus, genetic diversity mitigates the risk of pandemics – of course, that does not mean that pandemics cannot occur. Similarly, NISWIF may be able to mitigate large scale ($n \gg 1$) exploitation of software systems. The following discussion gives a basic mathematical proof for this hypothesis.

The cost of hacking (c_h) a system can be broken down into two components: a) cost of finding a vulnerability in the system (c_f), and b) cost of exploiting that vulnerability on that system (c_e). The cost of hacking one system can be written as:

$$c_{h_1} = c_f + c_e \tag{6.1}$$

Since the value of c_f is specific to a vulnerability, it is fixed for a vulnerability and it is different for different vulnerabilities. c_e quantifies the cost of engineering a malware that can successfully exploit the vulnerability on a given system. After a malware has succeeded in exploiting one system, it needs to be replicated to exploit additional number of systems. If c_r denotes the cost of producing a working replica of the malware to exploit a new system, the cost of hacking n systems can be expressed as:

$$\begin{gathered} c_{h_n} = c_f + c_e + c_r(n-1) \\ \text{for } n \gg 1, c_{h_n} = c_f + c_e + c_r n \end{gathered} \tag{6.2}$$

In the case of identical software systems, $c_r \approx 0$. Thus, equation 6.2 becomes:

$$c_{h_n} \approx c_f + c_e \tag{6.3}$$

From equations 6.1 and 6.3, we can express cost of hacking n identical software systems, c_h, as:

$$\begin{gathered} c_{h_1} \approx c_{h_n} = c_h \\ c_h = c_f + c_e \end{gathered} \tag{6.4}$$

It follows from equation 6.4 that given that systems are running on identical software the cost of hacking one system is approximately the same as the cost of hacking n systems. This implies that marginal cost of hacking every additional system by exploiting the same vulnerability must be zero. Taking the partial derivative with respect to n on equation 6.4, we can see that it is true:

$$\frac{\partial c_h}{\partial n} = 0 \tag{6.5}$$

Since marginal cost of hacking every additional system is zero, the adversary can manage to launch large scale cyber attacks without incurring additional costs. Employing a NISWIF strategy, we may be able to increase the marginal cost of hacking for the adversary. If the adversary incurs a reasonable marginal cost of hacking every additional system, then the costs of hacking large number of systems may become prohibitively high for the adversary to launch a large scale attack. To assess if this is true, let us begin with re-interpreting the equation 6.2 as follows using probability of successful exploitation p:

$$c_h = c_f + c_e + (1 - p)c_r n \tag{6.6}$$

For identical software systems, the probability that a replica of a previously successful malware may be able to successfully exploit a new system is ≈ 1. For non-identical software, if the method used to produce different variants of the software with identical functionality is strong, then we can safely assume that the probability that a replica of a previously successful malware may be able to successfully exploit a new system can be ≈ 0. Thus, using equation 6.6 we can explain the cost of hacking for both identical and non-identical software cases. Since the likelihood that a replica of a successful malware may succeed without any modification in exploiting a new system is quite negligible, it implies that the adversary incurs a reasonable cost of replication i.e. $c_r \neq 0$. In fact, one can assume that it is reasonably greater than 0. Thus, cost of hacking varies with the number of replicas to be produced, which in turn varies with the number of systems the adversary may want to exploit. Thus, taking partial derivatives with respect to n on both sides of 6.6 we can see that marginal cost of hacking in the case of NISWIF is:

$$\frac{\partial c_h}{\partial n} = (1-p)c_r \text{ where } c_r \neq 0, \text{ and } p \rightarrow 0 \tag{6.7}$$

From equations 6.5 and 6.7 we can summarize that marginal cost of hacking is:

$$\frac{\partial c_h}{\partial n} = \begin{cases} 0 & \text{for identical software} \\ (1-p)c_r > 0 & \text{for non-identical software} \end{cases}$$

This proves, at least theoretically, that a NISWIF strategy may be able to increase the adversary's cost of hacking and thereby help in mitigate large scale attacks. Note that a NISWIF strategy does not stop hacking altogether (i.e. $n = 0$), it only makes launching large scale attacks expensive (i.e. a '$n \gg 1$' problem may be reduced to '$n \geq 1$' problem). Such a reduction in the scale of the attacks can reasonably increase security of software systems. In the case of autonomous systems, for example, such an increase in security may translate to sizable economic benefits.

6.2 Is a NISWIF Strategy Feasible?

A NISWIF strategy appears to be technically feasible – it has been for quite some time now – but its operational feasibility has unclear. In fact, the lack of clarity on its operational feasibility may have been one of the reasons why not a lot of progress has been made on its technical implementation. This section briefly reviews the computer science literature that shoes it is technically feasible to develop NISWIF, and then we proceed to conceptualize a blockchain-based solution to demonstrate that it may be operationally feasible to distribute NISWIF.

6.2.1 Technically Feasible

The idea of securing systems through diversity of software has been around at least since the mid nineties (Forrest, Somayaji, and Ackley, 1997). However, it did not gain as much traction because of higher costs of development and maintenance, potential impacts on run-time efficiency, difficulty of tracking which variant of software was on which machine, and because Internet was expensive. Things are a little different today. Computation has become cheaper, Internet is more widespread, tracking versions of software has become easier and more efficient.

These changes have led to greater attention being paid to the idea of security through diversity. Since 2008, more than 25 academic papers have been published on this topic. Larsen et al. (2014) have done a great job of systematically surveying the topic of software diversity and critiquing the existing research in their summary of knowledge paper. While the research on software diversity is likely to increase over the next few years, it is fairly certain that software diversity is technically feasible. Most of its limitations originate due to operations and economic considerations more than technical considerations.

Software diversity techniques can be broadly divided into two categories: a) post-distribution diversity, and b) pre-distribution diversity techniques.

Post-Distribution Diversity

Post-distribution diversity techniques diversify (or randomize) the software at the client side, after the software has been delivered to the client machine, or after the software binaries have been created through the compilation of code. Some examples of post-distribution randomization techniques include address space layout randomization (ASLR), library entry point randomization, and binary rewriting. ASLR, for example, essentially adds randomness to locations of code loaded into the memory. If the adversary was exploiting the vulnerability of a particular module, then the adversary is likely to be interested in the relative location of that module in the machine's run-time memory. When ASLR is applied, the relative location of the module is not the same on every machine. This introduces some uncertainty, which affects the probability of successful exploitation. Other techniques use a similar approach as well, except the 'unit' that is to be randomized is different.

Pre-Distribution Diversity

Pre-distribution techniques, on the other hand, randomize the code that constitutes the software without affecting its intended functionality. Hence, it happens at the software manufacturer's side, before the code is compiled and assembled into binary packages. Some examples of this approach to randomization include reordering of instruction, randomization of function parameters, adding random but dummy branching-code inside functions, random reordering of functions, and system call mapping randomization. Essentially, all these methods randomize an 'unit of code' without affecting the functionality of the code. Each method has a different unit of code that it manipulates. A unit of code can be an individual instruction, a block of instructions, a function-call, a function itself, or the whole program.

Pros and Cons

While both approaches produce diversity in software, there are of course some pros and cons to both approaches. Some important pros of post-distribution methods are that they offer good software compatibility with third party software. Only minor changes are needed to the distribution mechanism. Diversification or randomization costs associated with these methods are quite low since a bulk of those costs are likely to be borne by the client. Cons of post-distribution methods are that they tend to be messy (or kludge-y) and difficult to implement. The number of transformations that post-distribution randomization methods can produce is limited. Lastly, these methods defend just okay against automated exploitation – low costs, limited benefits.

Pre-distribution methods, on the other hand, defend substantially well against automated exploitation. The scope for producing diverse variants with identical functionality is huge. Moreover, these methods are elegant and relatively easy to implement. The cons to this approach are: a) third-party software compatibility breaks down completely – every third-party software needs to be tailored for the instance, b) major changes are needed to the distribution mechanism, which make the operational feasibility of NISWIF rather unclear, and c) diversification costs are quite significant, which make its economic feasibility unclear.

Operationally Feasible through a Centralized App Store

Based on the above discussion, one can safely say that operating system manufacturers can produce NISWIF. However, distributing NISWIF is a challenge. One

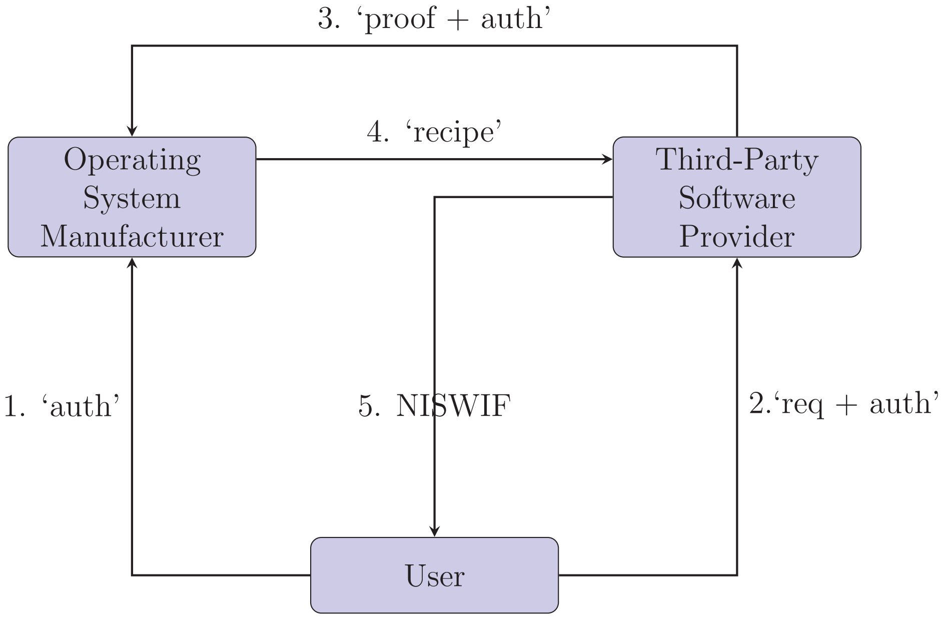

Figure 6.1: Operational Model for NISWIF in a Centralized App Store Context

way to address this challenge could be to use the currently popular centralized "App Store" model. The OS manufacturers can ship NISWIF to devices, and let third-party software providers to host their apps on the app store.

Whenever the user tries to install an app from the app store, the third-party app producer can request the recipe for randomization specific to the user from the OS manufacturer. However, to receive the recipe the third-party provider has to forward the authorization received from the user as the proof that it has been authorized to query the recipe. After verifying that the authorization provided by the third-party matches that received from the user, the OS manufacturer sends the user-specific recipe for randomization to the third-party. Upon being granted the recipe, the third-party producer can use it to produce the tailored app that is compatible with the randomized version of OS specific to the user.

Of course, this means that third-parties would need to get permission from the OS provider to host their apps on the app store. Such a permission-based model can ensure that adversaries may not be able to easily access the recipe for randomization specific to every user. However, even if the adversaries were to access the recipes they would still need to spend the computational effort to re-compile the malware according to the recipe for every system they intend to exploit i.e. their marginal cost would still be greater than 0. However, if adversaries find out a way to not bear those marginal costs or transfer those costs to systems that have already been exploited, then they can succeed in launching large scale ($n \gg 1$) attacks without incurring the costs associated with it despite NISWIF.

While a centralized solution of distributing NISWIF may be excellent for drones, UAVs, autonomous robots, and others that do not require a third-party app as such, it may create concerns of anti-competitive behavior in cases where third-party apps are needed. Since the OS manufacturer is the gate-keeper, all third-party producers have to seek permission from the OS manufacturer to deploy their apps on the app store *and* to deploy their app on every system that chooses to install it. The part of querying the OS manufacturer for a recipe every time a system wants to install their app is the one big difference between the proposed model and the existing app store models. In any case, the proposed model shows that a centralized model can feasibly distribute NISWIF to various systems (without the need for blockchain). However, the economic costs may be high for both OS and third-party providers. These costs are likely to be transferred to end users, especially in the case of free and open-source software.

6.2.2 Can Blockchain Help?

The above discussion shows that it is operationally feasible to distribute NISWIF through a centralized model. However, the centralized model (while perfectly well suited in cases where third-party apps are not needed) has some significant limitations. First, every operating system provider needs to have its own app store, which means every third-party provider has to seek permission from every operating system they want to develop their app for. This setting may raise concerns about anti-competitive behavior. Secondly, the economic costs of distributing NISWIF can be significant and are likely to be transferred to the end user. Third, there is no way to audit the transactions between the entities involved. Having access to transaction patterns may be quite useful in identifying the potential spread of a large scale exploitation effort, and a faster response is possible. Lastly, potential for beneficial network effects remains limited. A blockchain based solution may be able to address all these limitations.

Why Blockchain?

This section hypothesizes that blockchain may have the potential to bring software manufacturers, users, miners, and other important members to cooperate with each other under the framework of a blockchain based ecosystem that is economically conducive to a NISWIF strategy. Of course, there is a risk that a blockchain-based solution may be an overkill for the problem it is trying to address. However, there is at least one important benefit to exploring a blockchain-based solution to apolicy problem especially within the context of this dissertation.

Hypothesizing a blockchain-based solution to address a policy problem may give first hand insights into many of the complexities involved in designing such a solution and also into the new kinds of problems blockchain-based solutions may create. The lessons learned from this attempt at conceptualizing a solution may be transferable to other blockchain-based solutions being explored to address policy problems such as sharing medical records between healthcare providers, tracking land ownership rights, carrying out public elections and so on using blockchain. All of these problems are essentially the same class of problems. Each of them involves selectively sharing user-specific private information with entities who may legitimately need access to it, while protecting that information from falling into the wrong hands. The next chapter discusses the details of a blockchained ecosystem for distributing NISWIF.

6.3 Summary

In this chapter, we developed a theoretical proof as to how can non-identical software with identical functionality (NISWIF) can mitigate large scale ($n \gg 1$) cyber attacks by imposing a marginal cost to hacking. Next, we explored the technical feasibility of NISWIF and found that there is enough evidence to show that it is technically feasible. However, there are some operational and economic limitations associated with it. We also explored a hypothetical centralized app store model to address the operational limitations of NISWIF. While a centralized model can work excellently for systems that do not need third-party software to be install on the fly, it is not well suited for systems that need third-party software. Towards the end we hypothesized that a blockchain based solution may be useful in proving the operational and economic feasibility of NISWIF. The next chapter will explore this hypothetical blockchain based solution in greater detail.

7. A Blockchained Ecosystem

The problem with explaining a new idea instead of implementing it is: either the devil remains in the details or the explanation must be deviled with details

A complex network of individual entities interconnected with and interdependent on each other is called as an ecosystem. This chapter attempts to conceptualize a blockchain based ecosystem of software producers, users, miners, and other entities for the production and distribution of diversified software. The main goal of this chapter is to explore the operational and economic feasibility of producing and distributing software with the help of a blockchain.

7.1 The Ecosystem

The obvious entities in this ecosystem are: OS manufacturers, third-party software manufacturers, end users, and the nodes that form the permissioned blockchain. Since it is a permissioned blockchain, a central entity is needed to govern the blockchain. Because of concerns of anti-competitive practices, it is best for this central entity to be a non-profit organization representing the consortium of all major software producers. Facebook's Calibra blockchain has such an entity called the Libra Association – a group of companies have come together to form finalize the charter to govern the operations of this association and they will become the founding members of this association after the charter is completed.[1] The central

[1] See `https://libra.org/en-US/partners/` Accessed Jul-2019

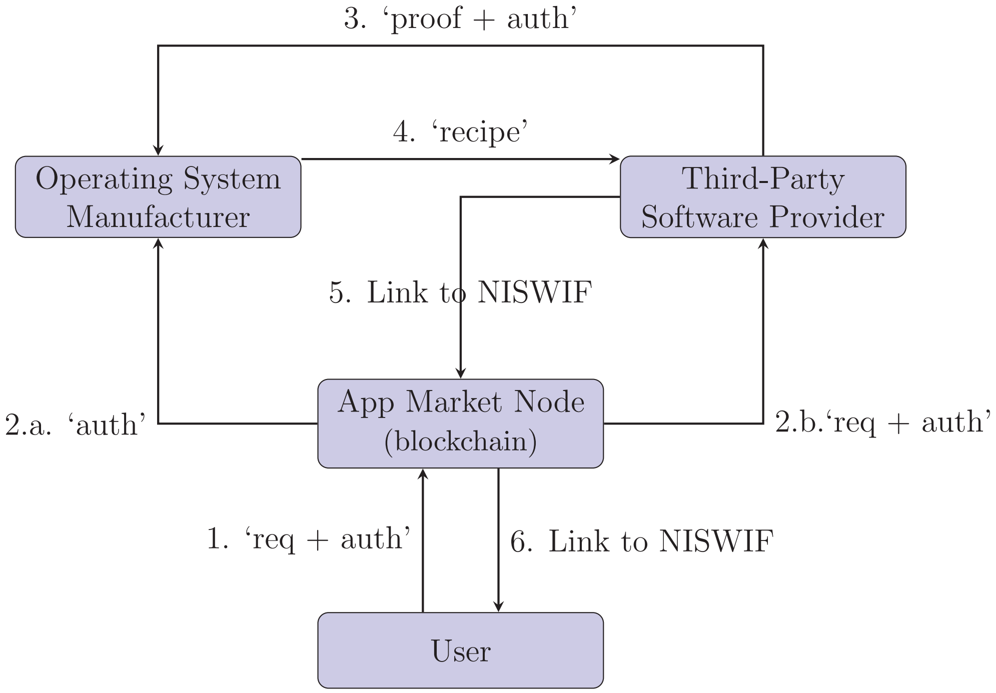

Figure 7.1: Distributing NISWIF in the Blockchain Ecosystem

entity acts as the gate-keeper that decides on who is allowed to be a member of this ecosystem, and what software are distributed through the ecosystem's decentralized app store .

7.1.1 Distributing NISWIF through a Decentralized App Store

Let us begin with clarifying the difference between a decentralized app and a decentralized app-store. By decentralized app (dApp) we mean an app that runs on a peer-to-peer network of computers as opposed to a centralized app that runs on a single computer. By decentralized app-store we mean that the app store itself runs in a decentralized manner on a peer-to-peer network. Through such an app store, a user can download a centralized apps (and decentralized apps too, if needed). The decentralized app store (or ‘app market’ from here on) acts as the trusted (but decentralized) intermediary that provides the user with user-specific diversified software.

As shown in Figure 7.1 (and discussed in detail in the Appendix), when the user chooses to install an app on her system, a message containing the request and the authorization is sent to the app market node (or a permitted node on the blockchain). The user's system only interacts with the app market when an app needs to be installed on it. Apart from this change, the rest of the distribution mechanism follows the same idea as discussed in the case of distributing NISWIF through a centralized app store.

However, one of the important things to note in a decentralized setting is that the user's private identifier, the operating system that the user's machine is running, and the recipe for randomization associated with that user should never be with any entity apart from the OS provider for reasons of security. If the adversary gains access to these thee crucial pieces of information, then the security of the user's system can be compromised. That said, the marginal cost of hacking argument still holds.

All the transactions between various entities related to production and distribution of software are recorded on the blockchain for reasons of transparency. Having such a log of transactions can provide useful insights into the origin and propagation of malware within the ecosystem. Of course, quasi-anonymity of identities must be maintained. It can be done using public-key cryptography techniques similar to those used in the bitcoin-blockchain, for example (more details in the appendix).

This shows that NISWIF can be distributed under a decentralized app market operational framework. However, the cost of compiling and assembling code into binaries can be prohibitively high. For example, if there are n users and the cost of compilation is x units, the total compilation costs can be at least nx units. This raises some severe cost concerns that may be enough to nip the idea of NISWIF in the bud. This is where blockchain may prove to be quite useful. The consortium can create a sharing-economy for computational resources using blockchain that may be able to alleviate the concerns about costs. The next section discusses such an economy in more detail.

7.2 A Blockchain Economy

There are at least two platforms available today for accessing shared computational resources on demand: a) Berkeley Open Infrastructure for Network Computing (or BOINC) developed by University of California, Berkeley,[2] and b) Golem blockchain.[3]

While BOINC is not a blockchain, it lets users (either individuals or organizations) sign up to provide their computational resources for research projects that require large scale scientific computations. Apparently, "[t]hese projects investigate diseases, study global warming, discover pulsars, and do many other types of scientific research." When a user signs up to be a volunteer node on the BOINC platform, BOINC software seamlessly downloads computing jobs to user's machines and runs them in the background. BOINC projected first started in 2002. In 2013, a cryptocurrency (and an associated blockchain) called GridCoin was created to provide probabilistic-reward or lottery based rewards to volunteers for providing their computational resources for BOINC projects.

Golem is another decentralized platform the supports renting of shared computing resources. The developers market it as "a decentralized network for sharing and leveraging large amounts of computing power." Golem uses an Ethereum-based payment system to remunerate computational resource providers with a crypto-token called Golem Network Token (GNT). Those is need for computational resources can rent nodes as computational resources and pay those nodes in GNT.

The underlying assumption of both the platforms mentioned above is that entities that need computational resources will seek out those resources on the Internet and

[2] See `https://boinc.berkeley.edu/` July-2019

[3] See `https://golem.network/` Accessed July-2019

pay the providers in their native tokens. This creates demand for those tokens, and the value of those tokens goes up since their supply is limited. This implies that value of these tokens comes directly from the actual and the speculated demand for them. This economic design principle is perhaps at the heart why these platforms have not gained a lot of popularity among computational resource providers.

When a user signs up to provide their computational resources to these platforms, they are essentially acting like Uber drivers giving rides to those who demand rides. Now these providers are expending their resources to complete the job i.e. they are spending their own money for powering the resource needed to complete the jobs they were hired for. However, since the value of these tokens appears to be completely driven by the demand for those tokens, they have proven to be highly volatile. If the user spent $\$x$ to complete the job, then the user would find it sustainable to continue providing computational resources only if the expected value of the token is greater or equal to the amount spent i.e. $E[token] \geq x$.

Since the tokens are not backed by anything, if their value falls below x the users may find it unsustainable to continue providing their computational resources to those platforms. Perhaps if these platforms did not force the transactions to be in their native tokens and let the transacting parties use a stable currency or even dollar-backed crypto-tokens, they may have a higher likelihood of being more popular. However, the problem in that case is that the platform would be subject to regulations such as KYC, and the platform would likely lose the quasi-anonymity features that cryptocurrency enthusiasts tend to hold in high regard. Not forcing transactions in their native crypto-tokens would also mean that the demand for those tokens would be quite limited. For tokens that derive their value completely based on the demand for them, such a limitation can be detrimental.

When designing a blockchain economy for shared resources where payments are made using a native crypto-token or a cryptocurrency, it is important to be aware of these problems. With this in mind, the next section tries to develop a simple financial model to show the cash flow and financial viability of sharing computational resources in the ecosystem described earlier in this chapter.

7.2.1 A Basic Financial Model for the Blockchain Economy

Let us assume a software licensing model for this economy i.e. users pay a certain fee per user per unit of time to gain a license to access the app market for diversified software. Let us say, for the sake of simplicity, that p_s is the average price paid by every user per unit of time (t) to gain one unit of license to access the app market. Note that a user may want to license $s_1 \dots s_k$ different units of software from this app market, and the license for each of them may be priced as $p_1 \dots p_k$ per unit of time. This would mean that every user pays:

$$\sum_{i=1}^{k} s_i p_i$$

to license k units of software per unit of time. If there are n users, then the total revenue per unit of time (R_s) from software licenses would be:

$$R_s = \sum_{n} \sum_{i=1}^{k_n} s_i p_i$$

Instead of carrying around this complexity of pricing software licenses throughout the model, we can use p_s to represent the average price of one unit of license to the app market per unit of time:

$$p_s = \frac{R_s}{n}$$

Let p_e be the average price of one unit of electricity needed to produce diversified software in this economy. Let p_l be the average price of one unit of labor needed for producing software in this economy. Let r_e and r_l be the average number of units of electricity and labor consumed respectively per license to the app market.

If q_s is the total number of licenses issued to the app market, then $q_e = r_e q_s$ will be the total number electricity units consumed (on average), and $q_l = r_l q_s$ will be the total number of labor units consumed (on average) to produce diversified software per unit of time in this economy.

Therefore, total revenue generated per unit of time in the economy would be, $R^{(t)} = p_s q_s$; total cost of electricity consumed per unit of time would be $C_e^{(t)} = p_e q_e$; and total cost of labor consumed per unit of time would be $C_l^{(t)} = p_l q_l$. For the sake of simplicity let us assume that there are no other costs are incurred in this economy and all costs can either be expressed in terms of C_e or C_l. Total profits generated per unit of time in this economy through software licensing can be expressed as:

$$P = R - (C_e + C_l) \tag{7.1}$$

Now let us introduce the ability of a user to share their computational resources to diversified software producers through a blockchain set up similar to BOINC or Golem. Software producers can outsource their computational jobs to users who sign up to be the nodes of the app market and the underlying blockchain. The nonprofit entity that manages the consortium's blockchain is the one that permits an entity to sign up as a mining node. The computational resources needed to produce diversified software are provided by the miners. This implies that the cost of electricity consumed to produce diversified software (mainly the additional compilation costs for diversifying the software) is borne by the miners instead of the software producers.

Thus, equation 7.1 can be expressed as:

$$P_{\text{sharing-economy}} = R - C_l \tag{7.2}$$

Note that the prices have not changed; just the cost of electricity is borne by the miners instead of the software producers. In return, miners are paid in the native token of this blockchain called *Eddy* ($\exists$).[4] Eddy is just an unit of account that measures the number of units of electricity spent for diversifying a software. Let c represent the conversion factor i.e. the number of units of electricity per unit of Eddy. Therefore, the value of 1 unit of Eddy is cp_e units in dollars. Since the nonprofit organization manages the blockchain for the consortium, it is the issuer of Eddy and is responsible for paying out all the outstanding units of Eddy. Users can redeem the value of the eddies they bear, in dollars, by submitting those eddies to the nonprofit organization. In other words, the nonprofit organization promises to pay the bearer of an Eddy the value that it represents, in dollars.

In order to pay the bearer of an Eddy the value it represents, the issuer must back those Eddies with an equivalent dollar amount. This amount comes from the difference between equations 7.1 and 7.2, which can be expressed as:

$$\Delta P = C_e$$

i.e. the costs saved, per unit of time, by the software producers is deposited (in dollars) at the consortium to back the value of outstanding units of eddies. Basically, by using blockchain as a ledger to track costs incurred by miners on behalf of the consortium and Eddy as an unit of account for those costs, the consortium of software producers can harness the computational resources of miners while freeing up some capital (C_e) in dollars that can be judiciously invested in financial markets. The returns on those investments (r_i) can appreciate (or depreciate) the value of eddies.

[4] The name Eddy is inspired by the concept of eddy currents in physics. Eddy currents are generated due to a changing magnetic flux of an electromagnet (i.e. an electric current induces a magnetic flux, which in turn induces an electric current). Similarly, transactions in the app market generate Eddies, which in turn may contribute to more transactions.

Therefore, the value of a unit of eddy can be expressed as

$$1Eddy = (1 + r_i)cp_e$$

Of course, since r_i is subject to market conditions it can be greater than, equal to, or less than 0.

In such a circular financial model, participating entities have some proverbial skin in the game and they can benefit from the network effects that a blockchain based economy can potentially generate. Acting in the interest of the consortium can generate benefits for individuals entities in the economy (and vice versa). Of course, acting in a way that is detrimental to the consortium will negatively impact individual entities too (and vice versa).

There may be something interesting to think about here. Politics, at least in western societies, has been fighting hard to bring capitalism into democracy; whereas blockchain appears to have the potential for bringing democracy into capitalism.

Anyway, since the entities that have eddies can use them for paying each other, it can be used as a currency for payment purposes within the economy. Secondly, as the value of eddies is backed by dollars, they have the potential to appreciate (or depreciate) in value based on market conditions. Third, having eddies enables an entity to participate in the economy even if they are not part of the consortium. These reasons may provide good incentives for bearers of eddies to hold on to them instead of immediately exchanging them for dollars upon receipt (or mining). These incentives to hold on to eddies coupled with the network effects from being part of the economy may contribute to some speculative demand for eddies, and by extension, may increase (or decrease) the value of eddies.

Again, as long as the expected value of an Eddy is greater than or equal to the cost spent for the number of units of electricity an Eddy represents (i.e. $E[Eddy] \geq cp_e$), the participants may find it beneficial to deal in eddies. The delta ($\Delta =$

$E[eddy] - cp_e$) comes from rate of return on investments and, from the demand for eddies inside (and outside) of this economy. Ideally, $\Delta \geq 0$ or $\frac{E[eddy]}{cp_e} \geq 1$ can indicate that eddies issued by the consortium are backed up by appropriate amount of dollars. That said, being aware of expected values is wise, and more detailed audits by regulatory agencies may be necessary to ensure that eddies are indeed backed up by dollars or other fiat currencies.

Limitations

While the above discussion may have provided some theoretical evidence that such a model is feasible, it is important to note down some of its important limitations. First, the challenge of designing a blockchained ecosystem appears to have become bigger than the problem of large scale cyber exploitation that it proposes to address. It introduces way too many moving parts – not all of which have been completely validated in the real world.

Second, prices of electricity may be different for different miners. Hence, using an expected value of price of electricity ($E[p_e]$) may be more appropriate. Expected values may sneakily introduce substantial miscalculations of prices and by extension, may show that eddies are completely backed up by fiat currencies when they are not (or vice versa). Also, price of electricity for commercial uses may be different than that for domestic/household uses. This may create mispricing of electricity and, by extension, misuse of electricity. If the scale of energy consumption is large enough (though it seems unlikely), it may affect the demand-load on electric grids.

Third, the nonprofit organization underpinning the consortium may end up operating like a bank under the proposed model. This implies that banking and other capital-related regulations may be needed to protect miners/investors from mismanagement of financial risks.

Lastly, as of now there does not appear to be any law or regulation in place that restricts a private entity from issuing a currency. While Eddy (like Facebook's Libra) is proposed to be backed by fiat currency, the consortium can – at a time when enough trust has been developed into their currency – decide to no longer back it with fiat by claiming that the value of their currency is completely dependent on people's belief in or demand for their currency.

Note that the above limitations (and their associated concerns) can probably be generalized to any blockchain based consortium economy that issues its own native token to be used as a currency in that economy.

7.3 Summary

This chapter showed that it is technically feasible to securely distribute diversified software using a decentralized set up. Such a decentralized set up when underpinned by a blockchain can enable sharing of electrical resources needed to produce diversified software. It also showed that if major software producers are willing to cooperate under a blockchain linked framework of operation, then it may be theoretically feasible to produce and distribute diversified software. However, there are some important limitations of such a set up that may make the challenge of setting the consortium up bigger than the problem the consortium is attempting to solve.

8. Concluding Thoughts

Many times strongly held
conclusions prevent us from looking
at something in a fresh new way

What I have done in every chapter of this dissertation is, I have followed a train of thought related to blockchain. Talk to anyone who does not know blockchain is, you will notice that if they are interested in that topic, they start following your train of thought. They will ask clarifying questions here and there but, by and large, they will follow your train of thought.

What I would like to do in this chapter is that I want to begin with building some context so that it is easier to bring the important trains of thoughts together without creating friction between them. This can potentially give the reader some useful insights into the phenomenon of blockchain and its implications without the noise that surrounds this topic. To ensure that the insights (or conclusions) are not taken out of context I have, deliberately, intertwined the context and insights with each other in this chapter.

Because I knew nothing about blockchain when I started this journey, I began with listening to (or reading) anyone who was somewhat coherently talking about blockchain. After a few days of listening to experts and reviewing the "spaghetti" literature on blockchain, a high-level pattern started emerging. The pattern was based on some of the trains of thoughts that I was able to catch and follow. I called them: a) the Implications, and b) the Applications trains of thought.

The Implications train was mostly filled with genuinely cautious questions like "what does this technology imply to the world we live in?" and "what problems is it likely to create for us?" and so on. On the other hand, the Applications train was filled with excitingly imaginative questions like "what can we do with this technology?" and "what new or innovative possibilities can we materialize using this technology?" and so on.

As one can imagine, there was a lot of friction between these two trains of thought because one was concerned with public safety (and security) while the other was concerned with entrepreneurial need for freedom to experiment. This friction contributed to a lot of noise that was drowning out the useful signals both the trains of thought had to offer. To separate the noise from the signal around blockchain, I had to structurally place those trains of thought in a way that did reduced the friction between them – otherwise it would defeat the purpose of this dissertation. Thus, I could not write any chapter without first harmonizing the conflict between those two trains of thought. Such deliberate structuring of thought enabled me to begin separating the signal from the noise.

8.1 Phenomenon of Blockchain

In the 2nd chapter of this dissertation, I attempted to follow one main train of thought: "what is blockchain?" I had to start with that question because many public policy researchers I spoke to who did not know what blockchain is, asked me that exact question. After going through many different definitions of blockchain, I felt I was stuck. Not many of those definitions were easy to comprehend. So I abandoned that train of thought temporarily and explored the question of "what does blockchain do?" This Verb frame of reference for looking at it provided a better structure to my thought than the Noun frame of reference I was looking through to understand what is blockchain. I began to comprehend what blockchain is by understanding what it does. The more I comprehended blockchain the more clearly I started observing the phenomenon of blockchain.

In simple words spread over the next few sentences, the phenomenon of blockchain appears to be the quest for creating a market where there is no necessity of a market maker for enabling transactions within that market. To create such a market, however, the inherent risks in the market must be appropriately realigned (to account for not needing a market maker anymore). Lastly, to attract participants into this quest of bootstrapping a market without market makers, value was created and provided as an incentive in the form of cryptocurrencies. This appears to be the core of the blockchain phenomenon.

8.1.1 The "Perfect Market" Hypothesis

It all started with Satoshi Nakamoto's paper that mysteriously appeared on the Internet in the immediate aftermath of the great financial crisis of 2008. Since nobody really knows what the end game of Satoshi Nakamoto's thinking was in developing blockchain, after spending a couple of years on this, I think I am in a position to offer an informed guess (to provide some context). The end game of Satoshi's thinking seems to be the creation of what I would like to call a "perfect market." I define a perfect market as one where the market maker is the market itself. In such a market, there is no one individual entity that is in a position to take systemic risks so large that may force the entire market into a crisis (if those risks materialized). It appears to me that, this is the logical conclusion blockchain may be headed (deliberately or randomly), if not Satoshi Nakamoto's train of thought. I will attempt to make it clear over the next few paragraphs why this may be the case.

Systemic risks are those risks that have a higher likelihood of forcing the entire market into a crisis. A market crisis is a market-wide gridlock of mismatches between buyers and sellers that brings a market to a halt. Similarly, an economic crisis is an economy-wide gridlock of mismatches between supply and demand. Broadly speaking, when entities that absorb risks on behalf of other entities in a market undertake risks of their own that are for whatever reasons not properly hedged, it exposes all the covered entities to the effects of those risks should they materialize. After the great financial crisis of 2008, it was clear that systemically vital market-making entities took on such systemic risks that, unfortunately, through unclear causes materialized into a financial market crisis (which inevitably resulted in a massive economic crisis).

It appears that market crises are conditional on a rather complex interplay between systemically vital entities and systemic risks, which is not easy to characterize. Furthermore, it is not clear whether it is a necessary condition or a sufficient condition (or both) either. The regulatory response to the financial crisis was to: a) thoroughly investigate the problem that led to the crisis, and b) based on those investigations enact new regulations that aim to curtail the legal ability of systemically vital institutions to undertake systemic risks (even if the intention is to spur rapid systemic growth). The regulatory view of these systemically vital institutions can be summarized as "too big to fail." This approach suggests that the regulators perceive the above mentioned condition as a necessary condition, if not sufficient as well.

On the other hand, Nakamoto's train of thought appears to be to exploring the following thought experiment: "we do not know if systemic risks can be eliminated. That said, is it possible to end the need for systemically vital institutions?" The logic being "if the need for systemically vital institutions did not exist then the complex interplay mentioned above could not exist either to precipitate a market crisis." Of course, this does not necessarily imply that market crisis could not occur at all. It is not unreasonable to expect that even the risk taken by a small entity may result in a market crisis through the mechanisms of a “butterfly effect”, which is a characteristic of chaotic systems like markets. However, what Nakamoto appears to be experimenting with is whether blockchain can spur a market where entities are “too immaterial to be vital” for such a market. Based on whatever is available, it appears that Nakamoto’s train of thought appears to be addressing the issue of “bigness” of entities rather than eliminating the possibilities of “failure” (or crisis) in a market. This suggests that just like the regulators’ train of thought, Satoshi’s train of thought, is interested in a systemically risk averse approach to economic growth, except through different means. It suggests that Satoshi is exploring “whether we can spur an economy that does not have opportunities for undertaking systemic risks?”

However, what this experiment appears to have missed is that many buyers and sellers appear to be wanting to avoid payment, settlement, or counterparty risks, even if it means relying on a market maker that may become too big to fail. For a small entity, individual risk of “failure” is more important to avoid than systemic risks probably because they tend to think at a local-level instead of a systemic-level or because it quashes their ambition to become big players in the market. The existence of such local risks can be perceived as a potential opportunity by a market making entity for adding value in the market by satisfying that need of an entity of not “failing” at clearing a transaction that may provide sustenance or growth (or both). Until the need for such addition of value is eliminated, the market opportunity for a market maker is likely to remain in tact. Since entities in financial markets treat money as a ‘store’ and an ‘unit’ of value, it (i.e. money) has loosely translated into the ‘source’ and ‘measure’ of value over time. The decision to create new value in an economy, however, tends to rest with one of the most systemically vital institutions called the central bank.

History shows us that, traditional economic train of thought has led us to a place where value has become the source of itself. However, not long ago gold was the source of value and fiat currency was the store of value, and the central bank was the legal arbiter between the source and stores of value. In many ways, it was probably inevitable that we would arrive at this stage where we had to equate (or conflate) the source of value with the store of value because of the issues and inefficiencies emanating from gold being a limited natural resource. The 'value' that gold enjoys is also the same reason for its rapid depletion on Earth over the past few centuries. Furthermore, the process of arbitration did not appear to be very efficient during emergency situations like wars or financial crises. Thus, it all kind of makes sense why we are at the place we are. However, this suggests that if the source of value has become the store of value, then adding value must ideally create new value into the system. It also suggests that if there was a ledger of value, there must be a match between total value created and total value accounted for in the economy because of this equation (or conflation) between source and store of value.

What the blockchain train of thought appears to be experimenting with is whether we can explicitly create a ledger of value where total value created and total value accounted for in the system. Satoshi appeared to be trying to achieve this by making blockchain the source of value and cryptocurrency the store of value. The process of arbitration between the source of value and the store of value is known as mining. The mining algorithms ensure that adding value would always match up with creating value in the system, and also, that this match is publicly verifiable even if the owner of that value has "lost" access to that value. Seen in this light, blockchain appears to be a digital metaphor for a natural resources based economy where the value is practically unlimited (2^{256}) but difficult to find.[1] To even begin finding this value, the miner (or the value hunter) needs a set of unconfirmed transactions. These unconfirmed transactions, in a way, provide the initial

[1] The number 2^{256} is an absurdly large number to fathom. It is *astronomically* larger than the total number of all gold atoms available on Earth.

directions or hints as to where the value may be available. The miners go in those directions to find value and add it to the system. Blockchain ensures that to create value, one must add value. The value that a miner adds in such a system is by clearing outstanding transactions. The transactions are confirmed only when there are no settlement risks. In the blockchain system, there is no lag between adding value to the market and creating value on the blockchain. Also, the entity that is adding value to the market is the same as the one creating value on the blockchain.

However, Satoshi's decision to put an upper limit (21 million bitcoins) to this practically unlimited value that could be created on blockchain is not very clear. That said, the purpose of this value is to serve as an incentive for interested miners to verify and validate the financial transactions occurring on blockchain. Because there is an upper limit to practically unlimited value, it can introduce a certain scarcity. Such scarcity perhaps makes the miners feel that it is worth spending their resources to find an opportunity to create and add value by verifying and validating unconfirmed transactions on the blockchain.[2] While this activity of verifying and validating transactions does address settlement risks (the risk that money does not get transferred at all), it leaves the users asking for more in terms of addressing payment risks (the risk that payment does not come on time). This time based thinking that underpins payment risk can make it difficult for an entity to engage in further transactions if that entity does not have other sources of money to cover that risk. It can even threaten the entity's survival in the market if they do not have enough money to engage in vital business transactions. Hence, at the local level in the market such entities feel the need to avoid payment risks too to mitigate the anxiety of going out of business if they are unable to secure (additional) source of funds.

[2] I did try to explore the relationship between value and worth in the beginning of the chapter titled 'Blockchain Created Value'. However, the topic itself felt like "sky is blue" kind of a topic to discuss.

This survival anxiety of an entity on the market has not been addressed very efficiently by blockchain's mining algorithms. These algorithms appear to be inefficient probably because of physical limitations of the medium on which blockchain operates. This inefficiency, as witnessed in the case of most blockchains, has led to a huge backlog of unconfirmed transactions for long periods of time. Due to this, processing of transactions becomes slower and the overall throughput of transactions is lower. It also shows that Satoshi's train of thought was more successful at systemically eliminating settlement risks but not payment risks. Given that his thinking appears to be risk averse at the systemic level rather than the local level in a market, this makes sense. But it also leaves the entities seeking more in terms of elimination of risks they face at the local level. Whether payment risks can actually be automated to be absorbed on blockchain remains to be seen.

That said, subsequent versions of blockchain like Ethereum did continue to experiment with eliminating payment and settlement risks through self executing smart contracts, and also, with eliminating counterparty risks through escrow mechanisms. There does not appear to be any momentum in favor of automating the absorption of payment risks as yet. It does appear that the payment risks may be a result of the physical limitations of medium on which blockchains operate. If payment risk gets automated in a way to be efficiently absorbed in a blockchain setting, it could bode well for the blockchain movement. Until then, it may just remain an interesting thought experiment for systemically risk averse thinkers like Satoshi (or Hayek for that matter).[3]

Furthermore, on the creation of value, the blockchain community's train of thought appears to be adamantly against conscientious but *ad hoc* manipulation of cryptocurrency units in circulation. The early blockchains appeared to be favoring a predetermined but completely self governing money supply algorithm. The

[3] I suspect that if Hayek was alive today, he would have been in support of carrying out such experiments.

‘predetermination’ aspect encoded in those algorithms determines the general time-schedule of money supply, and the ‘self governing’ aspect of those algorithms takes care of deciding exactly how many units must be entered into circulation at a given instant of time. However, the later generation of blockchains (like EOS) allowed for the option of conscientious but *ad hoc* manipulation of money supply algorithm itself through governance mechanisms such as voting and a community-drafted constitution.

The blockchain movement seems to be experimenting with eliminating the problems created by Satoshi’s unclear (or *ad hoc*) decision of putting an upper limit on the money that could be created on blockchain. This decision may also have enabled the decision of predetermining the supply schedule of money in the form of cryptocurrency units. One logic that can potentially explain Satoshi’s decision is that “if supply of money is *ad hoc*, then demand for money tends to be volatile. So, predetermined money supply schedules may reduce the volatility of demand for money.” However, wild swings in value of cryptocurrency units indicate that predetermined money supply schedules have counterintuitively transformed the problem of volatility in the supply of money into the volatility of demand for that money – all the volatility in the value of cryptocurrencies appears to be coming from the rising and falling speculative demand for units of those currencies. Thus, it does not appear to have succeeded in addressing the underlying problem of volatility that it may have hoped to address. Instead, it has transformed volatility from one side of money’s supply-demand equation to the other side. Perhaps, all of this was the result of the decision of setting an arbitrary upper bound on something that is practically unlimited; thereby, deliberately inducing demand that opens up opportunities for speculation.

8.1.2 Implications

Of course, the speculation fueled by such deliberately induced demand does not appear to be misplaced because the supply has already been predetermined. Any reasonably sound investor could be enticed by such an opportunity to make money through a pitch along the lines of "since money supply is fixed, the demand for it will be certain." In fact, that may have made early investors to rapidly adopt cryptocurrencies and advocate for them as well. This created an environment of hype that we all witnessed over the past few years. An environment where hype motivates investment instead of sound reasoning, is likely to be a breeding ground for fraud at its fringes. It sure was (and still is to some extent). It led to many of the public policy problems related to fraud that were discussed in the first part of this dissertation.

Upon deeper exploration of the problems that caught media's attention, I found that blockchain has not necessarily created new problems *per se*. However, it has the potential to scale up the intensity and the frequency of those problems if left unaddressed by law enforcement officials and regulators. It does not appear that new public policies need to be enacted to allay the collective fears expressed in the media about blockchain and cryptocurrencies.

That said, I would be contradicting Satoshi's train of thought on systemic risks if I did not point out that when the intensity and frequency of problems (of financial fraud, mainly) increase, it may lead to potential systemic problems that are too difficult to predict. If these problems become systemic enough in nature, then they can cause new problems that are more than legal nuisance in nature. Perhaps, they can even contribute to a market crisis of their own of some kind, if the stars are misaligned so to speak. Hence, the US SEC's (and other regulators' worldwide) seemingly hard-line regulatory responses towards new cryptocurrency startups do not appear to be unfounded. It suggests that if one agrees with what appears to be Satoshi's line of thinking on being systemically risk averse, then they have to agree with SEC's line of thinking too if they wish to be coherent in their opinions.

Cross-jurisdictional problems are likely to be one of the thorniest of issues that blockchain and cryptocurrency have exacerbated. The protection that criminals enjoy due to cross-jurisdictional issues between law-enforcement officials is also the reason why markets for illicit items have been so willing to accept payments in terms of cryptocurrencies. The major public policy gap that this trend has made apparent in the US (at least), is that of USPS being used as the default delivery mechanism for goods purchased on such markets. Since the USPS is a government agency, it is legally prohibited from intrusively inspecting the contents of a package addressed to a resident in the US. This loophole has ironically led a government agency to become the default delivery mechanism for illegal goods arriving from within and outside of the US. However, it was reassuring to find that postal intelligence officials are addressing this issue. Of course, it was understandable that not a lot of detail was available on this matter owing to the sensitive nature of these issues.

Another trend that appears to be burgeoning is that of creating "corporate currencies" using blockchain. Corporations appear to be grouping under a consortium to reap the benefits of network effects by creating a common corporate currency that can enable their customers to transact with each other within the economy created by that consortium. The most prominent manifestation of this trend is Facebook's ongoing attempt of creating a corporate currency called Libra. While such corporations are being extremely careful about abiding by existing public policies, it does appear that there is a public policy gap that may be exploited by the "ask for forgiveness instead of permission" line of thinking prevalent in the tech industry. Regardless of blockchain, in general, many prominent companies in the tech industry are adding a financial services arm to themselves. For example, Apple and Uber have issued their own credit cards; Google, Apple, and Samsung have their own payment mechanisms like Apple Pay, Google Pay, and Samsung Pay. The consortium of companies led by Facebook, is taking this trend to the next level through Libra. Today, Libra is a token (or a store of value) that is backed by a basket of fiat currencies of equivalent value. As such, it can be used as a currency in the corporate economy enabled by the consortium. However, there are no laws prohibiting the consortium from unbacking the Libra token if and when it becomes synonymous with currency i.e. the store of value becomes the source of value. This is a potential public policy issue that needs a lot more economic thought than what this concluding chapter can provide. Hence, I will leave it flagged here as a potential public policy issue.

That said, I will say that regardless of whether Libra (or any other corporate currency) succeeds or fails, it may be useful to explore the above mentioned potential public policy issue. There are many policy researchers who get stuck in the vortex created by the analytically resistant question of "whether technology should come first or public policy should come first?" when discussing these kinds of policy issues. What I learned while studying the public policy implications of blockchain is that it does not necessarily matter which one comes first as long as there is a match between them. I noticed that many of the public policies previously enacted much before blockchain was developed were robust enough to be able to address problems posed by blockchain. Now that technology is leading us to newer developments like corporate currencies, it may be a good opportunity for policy researchers to explore what policy responses can help us address any issues associated with such currencies.

8.1.3 Applications

The "blockchain based consortium" train of thought, however, appears to be something that technologists may do well to explore because of the potential for new technological innovations in this space, but also because it provides a for a newer way of collectively addressing a public policy issue created by an industry. One such policy issue is cyber security. The problem of cyber security is, essentially, a by product of the tech industry that no one player in the industry is able to coherently address. In the second part of this dissertation, I explored whether blockchain can be used to bring tech companies as a consortium to collectively mitigate the scale of cyber attacks.

Since one of the trains of thought in the first part of this dissertation was "systemic risks", I wanted to continue exploring that theme within the second part as well. It appeared to me that identical software creates a systemic risk to our collective cyber security. If a malicious entity can exploit one system, then that entity can likely exploit every other system running on identical software at a marginal cost that tends towards zero. This suggests that a cyber vulnerability that can be exploited at an individual level can also be exploited at a collective level with very little additional effort because of the homogeneity provided by identical software. This is one of the main reasons why we have seen many instances of comupter malware spreading rapidly around the world.

Seen in this light it becomes clear that the "identical-ness" of software is contributing to the systemic risk to our collective cyber security rather than the "vulnerability" of individual systems. This suggests that identical software is the mechanism through which individual risk can be turned into a systemic risk by a rogue or malicious entity. Breaking this mechanism may reduce the opportunities for any entity to threaten our collective cyber security. Breaking that link can reduce the scale of the problem from "$n \gg 1$" to "$n = 1$". With cyber-physical systems like autonomous vehicles waiting to gain regulatory approval, this issue becomes even more important to be explored for obvious reasons of national security.

As I explored the "identical software" train of thought further, it appeared that computer scientists have been exploring this problem on and off since the mid 90's but research on it has seen a definite resurgence since the late 2000s. I noticed that computer scientists are exploring if we can have machines with identical functionality without necessarily having identical software. Many of them are developing prototypes that demonstrate this capability to a certain extent. On the other hand, it was also clear that the software industry's economy is completely centered around the assumption of identical software. Companies develop software applications assuming the identical nature of an operating system on which the application will run. It is also one of the reasons why software companies are able to achieve the economies of scale that they are known to achieve. This suggested that even if computer scientists succeeded in developing non-identical software with identical functionality, it may not be adopted because of the way the software economy has built on the idea of identical software due to pragmatic reasons. This suggests that the challenge of safeguarding our collective cyber security cannot be addressed by one individual entity from the software industry who may produce non-identical software with identical functionality. It needs a collective response from the players in the software industry. It is possible that the "blockchain based consortium" model can find a significant role in spurring such an ecosystem that is centered around non-identical software to address the systemic risks created by identical software.

This made me explore the following question in the second part of my dissertation: "whether we can spur an ecosystem that is not centered around identical software using blockchain?" It ends with reasonably demonstrating that it may be feasible to use the "blockchain based consortium" model to mitigate a systemic risk to our collective cyber security. It sketches out the skeletal details of an ecosystem that is centered around non-identical software with identical functionality. It also conceptualizes a corporate currency called Eddy to balance the incentives of participating entities in such a consortium model. Of course, it is likely that such an ecosystem may have its own systemic risks. However, exploring those problems seemed to be beyond the scope of this dissertation.

8.2 Conclusion

Based on the above discussion, we can redefine blockchain as *a technology that can enable an economy where the market maker is the market itself.* It probably started out to be a technology to create an economy with a perfect market where the market maker is the market, and there is no player big enough to take risks that, if materialized, would force the market into a crisis. However, blockchain does not appear to have succeeded in its quest to produce a perfect market – whether it will succeed in this quest remains to be seen.

Since this is a dissertation in the domain of public policy, it may be appropriate to end it with some of the main policy recommendations. To begin with, I would recommend that allowing the blockchain experiment to continue going in search of a perfect market is not a bad idea. Given that it may be in search of an economy with a perfect market underpinning it, we may all be able to benefit from it if it succeeds in that quest. Since the scale and frequency of market crisis seem to be increasing, it is not a terrible idea to let the blockchain movement continue in its search for a market where no entity is systemically vital regardless of its size. A perfect market must naturally break the mechanism linking bigness of an entity to its systemic vitality to ensure no entity is in a position to force the market into a crisis. Thus, letting a technology that can help us experiment with realizing an economy with a perfect market can be beneficial on the whole despite the short-term problems it has created.

Also, I recommend that public policies that can ensure short-term problems do not contribute to systemic risks must be enacted as and when it becomes clear what problems can pose systemic risks. The current regulatory approach of regulating the "institutional interfaces" to blockchain appears to be much more effective as a strategy than regulating the development or deployment of blockchain itself. Secondly, it appears that identifying what problems can create systemic risks can be quite difficult. Even when such problems or vulnerabilities are identified, the system may be too centered on such vulnerabilities to do anything about it. In such cases, an industry wide response may be needed. Encouraging consortium based systemic solutions to address such risks can be a beneficial public policy.

Furthermore, to ensure that such problems are identified before they manifest themselves through systemic crisis, I recommend providing research grants to institutions in the field of complex systems and chaotic systems who are always exploring systemic risks. Better understanding and preparedness related to systemic risks can leave us better prepared to respond to systemic crises in the future.

Next, I recommend that policy makers encourage more experimentation with the decoupling of source and stores of value in an economy. The reason I recommend this potentially controversial view is because when the source of value and store of value are identical, the entities that possess the store of value can also easily get into a position of controlling that much of value in the system even if they do not have the legal right to create the stores of value they possess. That said, entities do have the legal right to create derivative stores of value that are not necessarily the sources of value in themselves like fiat currencies are. This is one reason economists have defined different money aggregates identified as M1, M2, and M3. Derivative stores of value allow for the total value accounted for in that economy to be mismatched with total value actually created by the central bank. By allowing for a mismatch between total value created and total value accounted for in the economy while having the source of value be the store of value, we may be enabling a significant source of systemic risk in our economy. Thus, decoupling the source of value and the store of value may ensure that even if the store of value is compromised during emergency situations, the total value in the economy still remains uncompromised. The degree to which governments may need to rely on law enforcement to enforce contracts during a crisis is likely to be much higher in the case where the source and stores of value are identical than in the case where they are decoupled. This is because possession of the store of value would not give an entity total control over the value stored in it, if they were decoupled. Hence, I recommend that we encourage the experimentation of decoupling the source and stores of value in our economies through the use of blockchain.

Furthermore, one of the biggest systemic risk to an economy emanates from national integrity because if national integrity breaks down, the economy as we know it would collapse too (or vice versa). In such a situation, having a central bank in charge of the store and source of value would not be likely to be of much use because it is likely to have lost its legal ability to regulate the economy. Just like the Internet was developed (partly at RAND) to ensure our ability to communicate with each other would remain in tact despite a nuclear attack, blockchain appears to be aiming at ensuring our ability to sustain an economy even if national integrity broke down due to an unforeseen emergency. Having the Internet along with an appropriate version of blockchain can help people reintegrate as a nation quickly if some random “once in a nation’s lifetime” kind of a systemic risk to national integrity materializes. A nation that can reorganize despite having briefly lost its national integrity, can be viewed as having one of the best national security policies a nation can hope for. If a nation wants to have a national reintegration plan as part of their national security policy, then promoting the experimentation of decoupling the source and store of value through blockchain may be important.

Thus, if I were advising the technology policy committee of such a nation that is thinking along the lines of a national reconstruction policy, I would recommend the promotion of experimentation with a National Blockchain. Whether having a national blockchain can safeguard the nation's capability to reengage in economic activities within itself and with other nations despite something like national disintegration has occurred? If a nation's capability to reengage in economic activities can resume quickly even after a catastrophic event (like national disintegration), then it must be a vital capability to protect. Such a capability can provide economic continuity to the nation or provide a sort of an "economic immortality" to the nation's economy. Providing economic immortality can be one of the most prominent applications of blockchain.

Appendix

Distributing Diversified Software

As shown in Figure 7.1, there are at least six important steps involved in the distribution of diversified software. This section gives some more details on how can this be achieved securely in a blockchain based decentralized app store setting.

Basic Cryptography

Public-Key Cryptography

It may be useful to begin with developing a basic understanding of public-key cryptography or asymmetric cryptography. In this cryptographic system, the sender can encrypt a message using the publicly available key of the receiver that can only be decrypted by the receiver using his/her private-key. For example, if the sender S wants to send a message M to the receiver R, then S encrypts M using R's public-key k_R i.e.

$$S \rightarrow R : \{M\}_{k_R}$$

Any message that has been encrypted using R's public-key k_R can only be decrypted by R. The public-keys are shared or easily available, the private-keys are kept secure and only the owner of the key can access it. This is why it is called asymmetric cryptography – one key encrypts, and the other can decrypt.

Symmetric-Key Cryptography

On the other hand, in a symmetric-key cryptographic system the same key k is used to encrypt and decrypt. This means that the key $\overleftrightarrow{k}$ has to be shared secretly between the sender (S) and the receiver (R).

$$S \rightarrow R : \{M\}_{\overleftrightarrow{k}_{SR}}$$

Protocol

The main aim of this protocol is to demonstrate the steps that can enable diversified software production in a decentralized app store setting based on blockchain. There may be other ways of orchestrating this protocol; the one discussed here is just one way of orchestrating it.

In this setup, no entity other than the trusted operating system manufacturer should be able to have a collection of user/machine specific id's along with the recipe for randomization associated with the instance of diversified software running on that machine.

The user (U) begins the process by requesting an App Store node (A) an app to be installed that is produced by a third-party software producer (T). The user's machine is running on an diversified operating system provided by Operating System manufacturer (O). U creates a request and an authentication and send it to A. Next, A forwards the request to O and the authentication to T. Next, T sends the authentication as proof that U has authorized T to request for the recipe for randomization associated with U's machine. Next, O verifies that the authorization provided by T is the same as the one that U sent to O through A. After the verification is successful, O sends the recipe to T. On receiving the recipe, T produces a custom version of software using the recipe. Upon producing the diversified software, T sends the download link to A. Lastly, A forwards the link to U who had requested for the diversified software.

Proof-of-work based nonces $(\eta_1 \ldots \eta_4)$ are used to maintain the integrity of messages, and to address delayed message-replay attacks. μ represents user-id; θ represents OS-id; τ represents third-party's id; r is the request and a is the authorization; δ_S represents the digital signature of the sender (S) of a message; R is the recipe for randomization; and lastly, l is the link to the diversified software.

$$\begin{array}{rll}
1. & U \rightarrow A: & \{\{\mu\}_{k_O}, \theta, \tau, r, a, \{r, a, \mu\}_{k_O}, \delta_U, \eta_1\}_{k_A} \\
2.1. & A \rightarrow O: & \{\{\mu\}_{k_O}, \tau, r, \{r, a, \mu\}_{k_O}, \delta_A, \eta_2\}_{k_O} \\
2.2. & A \rightarrow T: & \{\theta, a, \delta_A, \eta_3\}_{k_T} \\
3. & T \rightarrow O: & \{a, \delta_T, \eta_4\}_{k_O} \\
4. & O \rightarrow T: & \{R, \delta_O\}_{k_T} \\
5. & T \rightarrow A: & \{l\}_{k_A} \\
6. & A \rightarrow U: & \{l\}_{k_U}
\end{array}$$

The above protocol shows that at no point μ, R, and θ are in the possession of any entity other than the operating system manufacturer O. If those three are in the possession of an attacker, then the user's system can be compromised. However, the attacker has to still ensure that the malware is tailored to suit the user's machine.

Bibliography

Adler, John et al. "ASTRAEA: A Decentralized Blockchain Oracle". In: *arXiv preprint arXiv:1808.00528* (2018).

Ben-Sasson, Eli et al. "SNARKs for C: Verifying program executions succinctly and in zero knowledge". In: *Advances in Cryptology–CRYPTO 2013*. Springer, 2013, pp. 90–108.

Chang, Tao-Hung and Davor Svetinovic. "Improving bitcoin ownership identification using transaction patterns analysis". In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* (2018).

Cunliffe, Jack et al. "An island apart? Risks and prices in the Australian cryptomarket drug trade". In: *International Journal of Drug Policy* 50 (2017), pp. 64–73.

Décary-Hétu, David and Luca Giommoni. "Do police crackdowns disrupt drug cryptomarkets? A longitudinal analysis of the effects of Operation Onymous". In: *Crime, Law and Social Change* 67.1 (2017), pp. 55–75.

Dion-Schwarz, Cynthia, David Manheim, and Patrick B Johnston. *Terrorist Use of Cryptocurrencies: Technical and Organizational Barriers and Future Threats.* Rand Corporation, 2019.

Duffield, Evan and Daniel Diaz. *Dash: A privacy-centric crypto-currency*. 2014.

Ekblaw, Ariel et al. "A Case Study for Blockchain in Healthcare:"MedRec" prototype for electronic health records and medical research data". In: *Proceedings of IEEE Open & Big Data Conference*. Vol. 13. 2016, p. 13.

FinCEN. *FIN-2019-G001: Application of FinCEN's Regulations to Certain Business Models Involving Convertible Virtual Currencies.* 2019. URL: `https://www.fincen.gov/sites/default/files/2019-05/FinCEN%20Guidance%20CVC%20FINAL%20508.pdf`.

Fleder, Michael, Michael S Kester, and Sudeep Pillai. “Bitcoin transaction graph analysis”. In: *arXiv preprint arXiv:1502.01657* (2015).

Foley, Sean, Jonathan Karlsen, and Tālis J Putniņš. “Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?” In: (2018).

Forrest, Stephanie, Anil Somayaji, and David H Ackley. “Building diverse computer systems”. In: *Proceedings. The Sixth Workshop on Hot Topics in Operating Systems (Cat. No. 97TB100133)*. IEEE. 1997, pp. 67–72.

Hopwood, Daira et al. *Zcash protocol specification*. Tech. rep. Technical report, 2016–1.10. Zerocoin Electric Coin Company, 2016.

Irwin, Angela SM and George Milad. “The use of crypto-currencies in funding violent jihad”. In: *Journal of Money Laundering Control* 19.4 (2016), pp. 407–425.

King, Sunny and Scott Nadal. “Ppcoin: Peer-to-peer crypto-currency with proof-of-stake”. In: *self-published paper, August* 19 (2012).

Kumar, Amrit et al. “A traceability analysis of monero’s blockchain”. In: *European Symposium on Research in Computer Security*. Springer. 2017, pp. 153–173.

Ladegaard, Isak. “We know where you are, what you are doing and we will catch you: Testing deterrence theory in digital drug markets”. In: *The British Journal of Criminology* 58.2 (2017), pp. 414–433.

Larimer, Daniel. “Delegated proof-of-stake (dpos)”. In: *Bitshare whitepaper* (2014).

Larsen, Per et al. “SoK: Automated software diversity”. In: *2014 IEEE Symposium on Security and Privacy*. IEEE. 2014, pp. 276–291.

Mills, David C et al. “Distributed ledger technology in payments, clearing, and settlement”. In: *FEDS Working Paper No. 2016-095* (2016).

Mingxiao, Du et al. “A review on consensus algorithm of blockchain”. In: *Systems, Man, and Cybernetics (SMC), 2017 IEEE International Conference on*. IEEE. 2017, pp. 2567–2572.

Nakamoto, Satoshi. “Bitcoin: A peer-to-peer electronic cash system”. In: (2008).

Noether, Shen. "Ring SIgnature Confidential Transactions for Monero." In: *IACR Cryptology ePrint Archive* 2015 (2015), p. 1098.

Nofer, Michael et al. "Blockchain". In: *Business & Information Systems Engineering* 59.3 (2017), pp. 183–187.

Sankar, Lakshmi Siva, M Sindhu, and M Sethumadhavan. "Survey of consensus protocols on blockchain applications". In: *Advanced Computing and Communication Systems (ICACCS), 2017 4th International Conference on.* IEEE. 2017, pp. 1–5.

Sasson, Eli Ben et al. "Zerocash: Decentralized anonymous payments from bitcoin". In: *2014 IEEE Symposium on Security and Privacy (SP).* IEEE. 2014, pp. 459–474.

Sun Yin, Hao Hua et al. "Regulating Cryptocurrencies: A Supervised Machine Learning Approach to De-Anonymizing the Bitcoin Blockchain". In: *Journal of Management Information Systems* 36.1 (2019), pp. 37–73.

Swan, Melanie. *Blockchain: Blueprint for a new economy.* " O'Reilly Media, Inc.", 2015.

Tapscott, Alex and Don Tapscott. "How blockchain is changing finance". In: *Harvard Business Review* 1 (2017).

Walport, MGCSA. "Distributed ledger technology: Beyond blockchain". In: *UK Government Office for Science* (2016).

Zdanowicz, John S. "Trade-based money laundering and terrorist financing". In: *Review of law & economics* 5.2 (2009), pp. 855–878.