

2019

Knowing your bitcoin customer: A survey of bitcoin money laundering services and technical solutions for anti-money laundering compliance

Jesse Bryan Crawford
Iowa State University

Follow this and additional works at: <https://lib.dr.iastate.edu/etd>



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Crawford, Jesse Bryan, "Knowing your bitcoin customer: A survey of bitcoin money laundering services and technical solutions for anti-money laundering compliance" (2019). *Graduate Theses and Dissertations*. 17661.

<https://lib.dr.iastate.edu/etd/17661>

This Thesis is brought to you for free and open access by the Iowa State University Capstones, Theses and Dissertations at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

**Knowing your bitcoin customer: A survey of bitcoin money laundering services and
technical solutions for anti-money laundering compliance**

by

Jesse Bryan Crawford

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Information Assurance

Program of Study Committee:
Yong Guan, Major Professor
Doug Jacobson
Ying Cai

The student author, whose presentation of the scholarship herein was approved by the program of study committee, is solely responsible for the content of this thesis. The Graduate College will ensure this thesis is globally accessible and will not permit alterations after a degree is conferred.

Iowa State University

Ames, Iowa

2019

Copyright © Jesse Bryan Crawford, 2019. All rights reserved.

TABLE OF CONTENTS

	Page
LIST OF TABLES	v
LIST OF FIGURES	vi
ABSTRACT	viii
CHAPTER 1. OVERVIEW	1
CHAPTER 2. BITCOIN	5
2.1 Introduction	5
2.2 Cryptocurrency Concepts	6
2.2.1 Transactions	6
2.2.2 The Distributed Ledger	7
2.2.3 Blockchain	9
2.2.4 Proof of Work	14
2.2.5 Compensating Computation	18
2.3 Bitcoin	19
2.3.1 Bitcoin Standards	19
2.3.2 Bitcoin Transactions	22
2.3.3 Bitcoin Blocks	25
2.3.4 Reference Databases	27
2.4 User Experience	28
CHAPTER 3. ANTI MONEY LAUNDERING	31
3.1 Motivations and Methods	31
3.2 Money Laundering and Computer Crime	34
3.3 The Bank Secrecy Act	35
3.3.1 Reporting	35
3.3.2 Due Diligence	36
3.3.3 Anomaly Detection	37
3.3.4 Scope	38
3.3.5 AML Regulators	38
3.4 Know your Customer	39
3.4.1 KYC Principles	39
3.4.2 KYC Risk Factors	41

3.5	AML and Bitcoin	42
3.5.1	Applying Regulation to Cryptocurrency	42
3.5.2	Bitcoin Business Models	44
3.5.3	Enforcement	47
CHAPTER 4. BITCOIN LAUNDERING AND ANTI-LAUNDERING		48
4.1	Customer Identification and Mixing	48
4.2	Tumbling and Mixing In Practice	50
4.2.1	Centralized Mixers	51
4.2.2	CoinJoin and other multi-party transactions	53
4.2.3	Non-Tainting Mixers	54
4.3	Wallet Re-Identification in Practice	55
4.4	User Behavior and Bitcoin Implementation	58
CHAPTER 5. MIXING SERVICES AND AML SERVICES		60
5.1	Review of Bitcoin Mixing Services	60
5.1.1	Marketing	60
5.1.2	Features	61
5.1.3	Failure to Launch	62
5.1.4	Scams	63
5.1.5	Types of Mixers	64
5.1.6	Anonymity Protections	65
5.1.7	Operator and Law Enforcement Actions	66
5.2	Observations	67
5.2.1	Difficult to Establish	67
5.2.2	Community Response to Scams	67
5.3	Blockchain Analysis Services	68
5.3.1	AML/KYC Services	68
5.3.2	Investigative Tools	70
5.3.3	Implications for the Bitcoin Market	70
5.3.4	Dusting	71
CHAPTER 6. FUTURE DIRECTIONS		73
6.1	Improved Laundering Techniques	73
6.1.1	Multi-Signature Transactions	73
6.1.2	Output Splitting	75
6.1.3	Shuffling	76
6.1.4	Dining Cryptographers	76
6.1.5	Reliance on Anonymity Networks	77
6.1.6	Off-Blockchain	77
6.1.7	Altcoins	78
6.1.8	Observations	79
6.2	Bitcoin Reidentification	79

CHAPTER 7. CONCLUSION	81
7.1 Bitcoin Laundering	81
7.2 Bitcoin Anti-Money Laundering	81
7.3 Recommendations for Mixing Services	82
7.4 Recommendations for Bitcoin Services	83
7.5 Recommendations for AML Regulation	84
7.6 Fundamental Conflict	85
BIBLIOGRAPHY	87
APPENDIX A. LIST OF MIXING SERVICES	87
APPENDIX B. LIST OF ANALYSIS SERVICES	91

LIST OF TABLES

	Page
Table 2.1 Evaluating a P2PKH Script	20
Table A.1 Mixing services evaluated	57
Table B.1 AML risk scoring services	61

LIST OF FIGURES

	Page	
Figure 2.1	Structure of example cryptocurrency transactions. Transactions can have an arbitrary number of inputs and outputs, which may “come from” and “go to” any other transaction, forming a graph.	6
Figure 2.2	Structure of example cryptocurrency blockchain. Each block refers to the hash (actually Merkle tree root in the Bitcoin implementation) of the previous block.	9
Figure 2.3	Structure of the blockchain, showing a fork which has been abandoned by miners after another fork overtook it in length, ensuring a single longest chain.	12

GLOSSARY

AML: Anti-money laundering regulations.

BSA: The Bank Secrecy Act, which introduced AML requirements to the banking industry.

CIP: Customer identification program, a core component of KYC compliance.

FinCEN: The Financial Crimes Enforcement Network, a bureau of the US Treasury Department.

KYC: Know your customer, a typical component of an AML compliance program which requires financial institutions to identify their customers and establish bounds for normal behavior.

P2PK: Pay-to-public-key, the most common type of Bitcoin transaction in which the spender of an output must provide a signature from the public key the output was addressed to.

P2SH: Pay-to-script-hash, a more recently introduced type of Bitcoin transaction where alternate rules are used in verifying spending of an output. Instead of the output being addressed to a public key, it is addressed to the hash of an arbitrary Bitcoin script which is provided by the spender.

UTXO: Unspent transaction pool, the set of Bitcoin outputs which have not yet been used as inputs to other transactions.

ABSTRACT

Cryptocurrencies are gaining significant attention and financial investment. Among the wave of new cryptocurrencies, the first cryptocurrency introduced, Bitcoin, remains the most notable and most heavily used.

While Bitcoin is often perceived as an anonymous system, it is in fact only pseudonymous and a variety of methods are known to reidentify the holders of Bitcoin wallets. As a result, services have emerged which “anonymize” Bitcoin by making it difficult to trace the origin of Bitcoin funds. These services are referred to as “mixers” or “tumblers,” but are more generally methods of laundering Bitcoin funds.

In the United States, a system of anti-money-laundering (AML) regulations developed since the 1970s requires financial services organizations to take positive steps to identify their customers, prevent use of their services for money laundering, and detect and report customers which appear to be engaged in money laundering.

These AML regulations have been interpreted by the primary regulator, FinCEN, as fully applicable to Bitcoin. This creates a clear conflict with laundering services which are directly intended to prevent organizations identifying the possessor of funds.

This thesis explores the advancing state of both Bitcoin laundering services and Bitcoin anti-laundering services intended to assist in compliance with AML regulations. The current state of the art in both laundering and anti-laundering services is explored. Later, current research and avenues for improvement in these services are discussed.

Ultimately, the way forward for Bitcoin AML regulation is discussed. The current regulatory approach to Bitcoin is insufficient to mitigate laundering with Bitcoin and should be refocused.

APPENDIX 1. OVERVIEW

Since its introduction in 2008, by a whitepaper from an apparently pseudonymous author, Bitcoin has attracted a great deal of attention and investment. The Bitcoin community is large and vocal, and Bitcoin has ushered in a larger family of cryptocurrencies and other applications of blockchain technology.

One of the most significant applications of Bitcoin is in privacy, anonymity, and circumvention of regulation. However, Bitcoin is by nature a “distributed ledger”, meaning that the ledger of transactions is readily available to anyone who wishes to inspect it. The ledger is in the form of a graph in which addresses are nodes and transactions edges, and so is amenable to network analysis techniques. This challenges the anonymity features of Bitcoin, since a wide variety of methods can be used to re-identify Bitcoin users based on their recorded behavior.

The demand for privacy despite this limitation has led to services referred to as “mixers” or “tumblers”, which are the Bitcoin equivalent of money laundering. These services create artificial transactions which, due to the nature of the Bitcoin system, can conflate or “mix” units of Bitcoin in such a way that it is difficult or impossible to determine which coins originated where.

These services are useful to Bitcoin users because they break the chain of possession for coins. This may be important in cases where, for example, Bitcoin was obtained as a result of illegal activity or in payment for a service or good provided anonymously. After mixing, the owner of such coins can spend them without others tracing their origin.

These mixing services are an interesting topic for study in two ways. First, they present an interesting application of cryptographic and stochastic methods to prevent algorithmic analysis of the transaction graph. While clearly useful in the commission of crime, these methods also have implications for a variety of computer and information privacy challenges. Second, an understanding of Bitcoin mixing systems in use will be critical to law enforcement and financial regulators in detecting and prosecuting the increasing number of crimes committed using Bitcoin. The operation of these mixing services in practice, including some programmatic analysis of their behavior, will be discussed at length.

Bitcoin mixing services have arisen at the same time that the financial regulatory system in the United States has worked to apply regulations to cryptocurrency. One of the most important and complex components of US financial regulation is anti-money-laundering or AML regulations. These impose certain requirements on financial institutions to deter and detect money laundering by their customers, which include a set of expectations referred to as “know your customer” or KYC. KYC regulations mean that financial services are expected to perform due-diligence monitoring of their customers activities and background. Businesses that offer financial services in Bitcoin, including trading Bitcoin for traditional currencies, are now generally required to comply with these regulations.

AML and KYC are difficult to apply to Bitcoin, both because the Bitcoin system is designed to provide pseudonymity over traceability and because Bitcoin is often involved in complex transactions dissimilar to those seen with traditional currencies. However, the open nature of the distributed ledger allows for the use of network analysis techniques to detect whether or not funds deposited by a customer are the result of laundering. Additionally, the very nature of Bitcoin

calls into question the practicality of refusing laundered coins, since over time a large portion of total Bitcoin currency will have some history of laundering. This work discusses the application of United States anti-money laundering regulation to Bitcoin in detail, including the several rounds of regulatory guidance issued by FinCEN.

Unlike existing publications in this area, which focus primarily on theoretical designs or on a small selection of real laundering techniques, this work attempts a comprehensive review of the landscape of laundering and anti-laundering services which exist now or had previously existed.

To investigate the broader ecosystem of Bitcoin laundering, this work includes an effort to enumerate all current and historic Bitcoin laundering services including basic information on their mechanism of operation and fate, when possible. This effort includes investigation of the ways that these mixing services advertise and describe themselves, and the ways that users select mixing services for use. The marketplace of Bitcoin laundering services is complex due to the risk users face of simply losing their funds to fraudulent services, and so an informal system has formed to monitor mixing services for trustworthy behavior.

While more advanced Bitcoin laundering methods based on multi-signature transactions can reduce the risk of fraud and poor anonymity due to design defects, such multi-signature mixing systems are uncommon in actual usage. The majority of Bitcoin laundering is conducted using conventional centralized mixing services in which the laundered coins are entrusted to the service's operator. This likely reflects the improved usability and privacy aspects of these services.

While there has been one high-profile incident of law enforcement seizure of a mixing service, outside interference with mixing services is uncommon. The vast majority of mixing services fall to a much simpler fate, a simple lack of interest from the community. Mistrust of new laundering services makes it difficult for them to get off the ground, and this mistrust is apparently founded as many laundering services appear to be scams or otherwise fraudulent.

The need to investigate criminal activity and achieve AML compliance for Bitcoin services has led to a number of tools and services directed towards detecting and investigating Bitcoin laundering. Developers of these tools publish very little about their methods, likely out of concern that they could be defeated by knowledgeable launderers.

Most laundering analysis tools are intended for manual use and are largely limited to visualization of transaction relationships and tagging and tainting of transactions. However, a new class of services have arisen which provide automated risk scoring of wallet addresses intended to meet AML compliance requirements. These services are of particular interest since they must attempt to detect laundering automatically. Operators of these services do not disclose their principals of operation, but they are likely based on taint analysis from addresses known to be involved in laundering services.

One prominent laundering service has attempted to interfere with taint-based detection of users who have laundered Bitcoin by sending small, unsolicited deposits to a large number of Bitcoin users, thus tainting their wallets with laundered funds. While interesting, this "dusting" attack is unlikely to be effective, since it is easy for mixing services to detect and ignore this behavior.

Next, this work discusses the current frontier of development in Bitcoin laundering services. While it is possible to design Bitcoin money laundering systems which create perfect anonymity within a set of mixing partners, and these methods have been known for some time, practical limitations mean that these methods are infrequently used. Application of a user-centric lens to these methods reveals their undesirable properties and poor ease of use. Methods of Bitcoin laundering in the academic literature are generally impractical because they require mixing of only

fixed amounts (reducing the possible anonymity set) and have other properties which reduce the size of the anonymity set.

Ultimately, the profit motive of most Bitcoin laundering services is a further deterrent to the use of multi-signature methods, since these decentralized services lack a clear method of monetization.

Finally, this work includes some recommendations for the developers and operators of Bitcoin laundering and anti-laundering services. It is most likely that Bitcoin laundering will continue to be done primarily using centralized services for the foreseeable future, but there are opportunities to improve centralized mixing services and to design decentralized, multi-signature mixing systems which are more attractive to users.

Services and tools for analysis of Bitcoin laundering, while difficult to discuss accurately due to their closed nature, appear to rely on relatively simple techniques and may be vulnerable to manipulation.

The current regulatory approach to cryptocurrency has been developed ad-hoc by extending existing money transmitter regulations to cryptocurrency businesses and enforcing AML requirements “at the edge” when traditional currencies are exchanged to or from cryptocurrency. This method is insufficiently comprehensive and leaves many opportunities for unobserved laundering activity. To be effective in the AML mission, regulation must be developed for cryptocurrencies themselves, accounting for the unique properties of cryptocurrency systems.

APPENDIX 2. BITCOIN

2.1 Introduction

Bitcoin, and cryptocurrencies more generally, are proposed for a large range of use cases. In theory, Bitcoin can be used as anything from a long-term store of value (due to the security of its cryptographic protections) to a medium for micropayments (due to the low cost of recording transactions). Since the introduction of Bitcoin a large number of cryptocurrencies (often referred to as “altcoins”, this being a portmanteau of “alternate” and “bitcoin”) have been introduced. Many of these cryptocurrencies are very similar, while some form “families” with significant differences from bitcoin. The overall landscape of cryptocurrencies will be discussed in depth.

Bitcoin is, in different ways, suitable as a replacement for cash, electronic payments networks (such as those which process credit cards), and long-term savings vehicles. Bitcoin’s characteristics differ from existing payment systems in such a way that it can be difficult to understand in comparison to them, and so rather than discussing Bitcoin in contrast to existing payments systems, its technical architecture will be outlined followed by a discussion of each of its common use-cases.

While the potential applications of Bitcoin are broad, in practice, Bitcoin has certain technical advantages which make it attractive to certain users. Bitcoin’s decentralized nature is appealing to individuals who distrust governments or banking authorities, while its pseudonymous nature is advantageous to those who wish to complete financial transactions anonymously. This has made Bitcoin particularly common in illegal and gray-market commerce. On the other hand, Bitcoin’s recent introduction, popularity, and cryptographic security also make it popular as a speculative investment.

The pseudonymous nature of Bitcoin and the complex transactions found in use make it difficult to determine typical usage patterns for Bitcoin. As a result, most assertions about “the primary use of Bitcoin” are speculative. Efforts at rigorously determining Bitcoin’s popularity and applications will be discussed.

2.2 Cryptocurrency Concepts

2.2.1 Transactions

Cryptocurrencies are built on transactions: they track the way that currency flows from one user to another. A sometimes counter-intuitive implication of this design is that cryptocurrencies are not particularly concerned with the balances of accounts. The amount of currency that an individual possesses simply falls out of the differences in transactions that they have received and the transactions they have sent.

Unlike conventional currencies, there is no physical token or central authority to confirm that a person is entitled to spend the money in a transaction, and so some mechanism to verify that transactions are valid and authorized is required. Cryptocurrencies verify transactions based on cryptographic signatures. As a result, a “wallet” or “account” in a cryptocurrency system is defined by possession of a cryptographic key.

Generally, each transaction consists of metadata (e.g. a timestamp), inputs, outputs, and a signature. The inputs on the transaction refer to previous transactions that are being "spent", while the outputs typically indicate the key of a recipient, but more broadly specify the conditions under which the transaction can be spent by future transaction. The transaction is signed by the key pair that the input transactions specified, indicating authorization to spend those coins.

Consider an example: if you have previously been paid 10 coins and 5 coins, in terms of the cryptocurrency system this actually means that there are transactions specifying an output of 10 and 5 coins respectively, and whose outputs require a signature by your key pair. You can then send 15 coins to another person by creating a transaction which refers to the previous two transactions as inputs, provides a signature verifying your authorization to spend those previous transactions, and specifies an output that requires a signature by the key belonging to your intended recipient.

The outputs are actually significantly more complicated than the address of a recipient. Bitcoin introduced the concept of an *output script*, and because of the flexibility this provides it has been broadly adopted in other cryptocurrencies. An output script is, in actuality, an arbitrary program written against a simple stack-based virtual machine. A set of operators are provided, including cryptographic operators, that allow for the construction of complex criteria to spend a transaction that may include multiple separate signatures.

Like outputs, the inputs of a transaction are also more complex than they may seem. In many cryptocurrencies the input of a transaction can also include a script, allowing particularly complex forms of transactions.

The number of inputs and outputs on a transaction are arbitrary and may sometimes be large. The ability to include multiple outputs is particularly important because of a simplifying property of most cryptocurrency systems: a transaction must spend the entirety of its inputs. Because it is often not possible to sum inputs to the desired output, it is common for transactions to include at least two outputs, of which one is the *change output* which returns the excess balance of the inputs back to the spender. For example, if you wish you to send a person five bitcoin but only have a transaction of ten bitcoin unspent, you will generate a transaction for ten bitcoin, of which five go to the recipient and five back to yourself.

It may seem that cryptocurrency transactions are complicated, and indeed, they only become more complicated when a specific cryptocurrency is examined in depth. The potential complexity of transactions is one of the most challenging aspects of cryptocurrency analysis.

2.2.2 The Distributed Ledger

The wider field of cryptocurrency was, for most purposes, conceived by the pseudonymous paper and reference implementation introducing the Bitcoin system in (?). Most cryptocurrencies today are derived from the Bitcoin design and share a great deal in common with it. The most fundamental concept, which is basic to all cryptocurrencies, is that of the "distributed ledger." This system is most easily explained in light of the most fundamental problem that currency systems face: a person spending more than they have.

In conventional currency systems, such as fiat paper currency, this is prevented by the use of physical tokens. If a person gives you ten dollars in cash, you have a strong assurance that the money is available to the spender by simple merit of them having possessed it in order to give it to you. You need not determine the origin or provenance of that money (although in some cases this may, in fact, be advisable, something which will be discussed later).

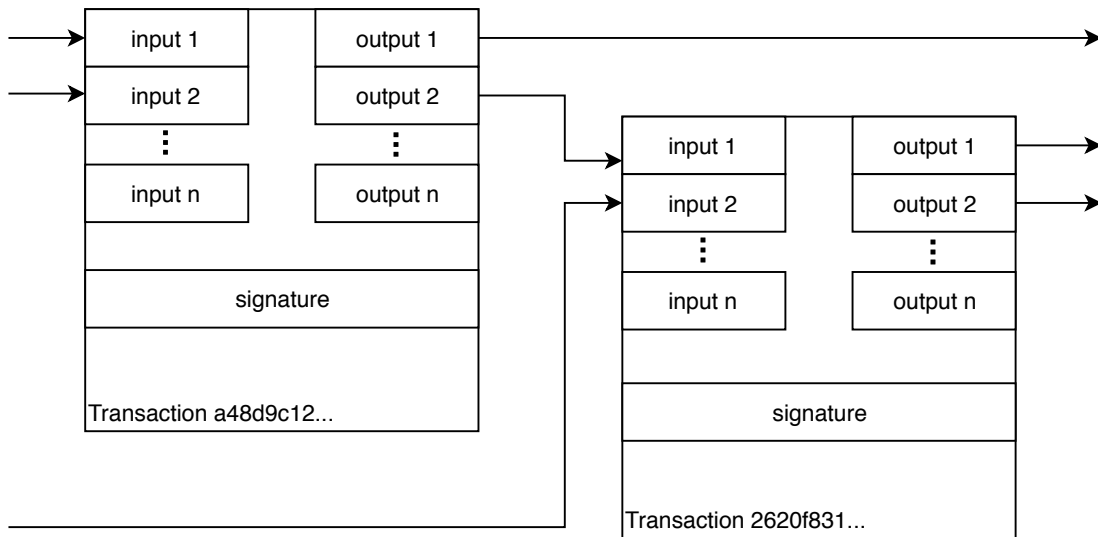


Figure 2.1 Structure of example cryptocurrency transactions. Transactions can have an arbitrary number of inputs and outputs, which may “come from” and “go to” any other transaction, forming a graph.

This same model is not effective in the case of electronic currency for the simple reason that the data representing an amount of currency can trivially be duplicated. This permits a behavior called “double-spending:” if a person gives you ten electronic dollars, it is quite possible that they have given the exact same currency to ten other people, a trivial form of fraud that is so simple and effective as to yield the currency system useless. So, it is necessary that there be some system to ensure that a person does not spend more money than they have. An obvious solution to this problem is a ledger.

A ledger system will be familiar to most people, as it is the system used by banks and electronic payment networks to track their customer transactions. If a trustworthy central authority (such as a bank) were to keep track of the amount of money a person had, you could then verify with that authority that a person giving you money has it available to spend. The transaction can then occur without any physical tokens, by the spender’s bank agreeing to decrement their balance and your bank agreeing to increment yours. The vast majority of money in the financial system today moves in this fashion, with various types of payment networks facilitating verifying a person’s available funds and moving those funds between institutions.

This system is workable but has a clear defect: it is reliant on trustworthy central authorities. Bitcoin was developed with the goal of addressing this problem by keeping a *distributed* ledger instead of a centralized one. The distributed ledger is available and readable to everyone participating in the network. If a person pays you ten Bitcoin, you can verify in the ledger that they actually possess those ten Bitcoin. Once the payment is made, you can verify in the ledger that *you* now possess those ten Bitcoin and are likewise permitted to spend them.

The ledger must be open to all and auditable for correctness. That is, all users of the system must be able to verify that a given transaction is legitimate (the sender has enough funds and the transaction has been recorded as going to the recipient), and all users must be able to verify that

the copy of the ledger they are inspecting is the one true and correct version, will be seen by all other users of the system, and will not be changed.

Distribution of the ledger to every participant in the network is a readily solvable problem using existing peer to peer network technology. Auditing the ledger for correctness is a more difficult problem, one which Bitcoin introduced a novel solution for: the blockchain.

2.2.3 Blockchain

Although the blockchain concept introduced in [1] was largely considered novel at the time of its release, there is prior art for the blockchain concept¹, which is helpful in explaining the principal. [1] describes, and [2] refines, an approach to time stamping documents to prove their existence at a given date. This is a useful ability in the areas of e.g. intellectual property law. Consider the time stamping of documents as an example case in which, like a public ledger, the nature (documents) and ordering (time) of events must be verifiable.

2.2.3.1 Auditable Time Stamping

Using well-understood cryptographic methods, an author can generate a cryptographic hash of a document (using e.g. a SHA family function) and submit it to a time stamping authority (TSA). The TSA can then append the timestamp and cryptographically sign the result, producing a verifiable assertion by the TSA of the time that the document (or rather, its hash) was received.

This approach has the significant downside of depending on the honesty of the TSA. If the TSA were to create a false timestamp, it would still appear genuine. To prevent dishonesty by the TSA there must be some mechanism to audit the issued timestamps for correctness. [2] introduces two general approaches to this problem, which were more fully developed in further work by the same authors including the establishment of the company Surety which implemented the concept [2].

The first solution is referred to as *catenate certificates* in [2] although the term *linked time stamping* is more common today. A TSA providing a public service will receive requests to timestamp and sign document hashes on an ongoing basis, and will generally not have any control over the contents or ordering of these requests. This can be exploited to make improper timestamps detectable. If each signed attestation generated by the TSA included not only the document hash and the timestamp, but *also* an artifact of the most recent timestamp previously issued by the TSA (e.g. a cryptographic hash of the attestation), the linking of one attestation to the next creates a *chain* which allows verification of the ordering of the timestamps.

If the TSA were to issue a false timestamp, an audit conducted after the fact would find that the false timestamp were mysteriously missing from the point in the chain where it should have appeared (between the first document before the false timestamp and the first document after the timestamp). In fact, it may appear in the chain in another location, or may be missing entirely. This mechanism enforces that the TSA be truthful about at least the order in which documents were timestamped, and cannot be defeated unless the TSA is able to predict the hash (and therefore contents) of the falsely timestamped document at the time that the first legitimate timestamp after is generated, which limits the attack to being nearly useless (since the contents of the document falsely timestamped must be known—that is, the document must exist—as of the time of the actual timestamp).

¹Nakamoto was aware of and, in fact, cited this work, which has nonetheless been often overlooked as an important contribution to the present field of cryptocurrency

Alternately, the TSA could produce a false timestamp even under a linked time stamping scheme by simply falsifying the entire chain following the false timestamp. This is not detectable after the fact. Still, if the volume of documents being timestamped is high and the falsification is over a relatively short period of time, this may be practical. Preventing it entirely requires a mechanism to ensure that any auditor is aware of a single valid chain and can disregard any "false" chain fraudulently produced to legitimize a false timestamp.

This second approach, implemented by the company Surety, adds an additional layer of security. Even if the chain has branches or forks (as a result of fraud or error), a single "tip" of the chain can be established as the legitimate chain if it is made publicly known. Once a chain has been publicly known as the legitimate chain, any new chain can be ignored as meaningless.

First, a tip of the chain must be selected as legitimate. In the case of Surety, the legitimate chain was simply selected by the operating company. Requiring that the selected chain be based off of the previously selected chain (that is, the legitimate chain must be an extension of the previous legitimate chain) and that the legitimate chain be published at frequent intervals mitigates this naive selection process since it would only allow a fraudulent chain to be created back to the most recent previous publication.

Second, a mechanism must be found to publicize the legitimate chain. Surety used a cryptographic summation method in which a hash was generated which covered every timestamp in the chain since the previous publication. The choice to include every transaction in the chain, rather than only the most recent, allows this to serve as an additional attestation of each individual transaction as well, guarding against any vulnerability in the chaining mechanism. The cumulative hash can be generated by one of a number of methods, typically using a sequential or tree structure. This hash was then published as a classified advertisement in the New York Times [?].

In summation, a timestamp can be audited for legitimacy by the following steps: First, the most recent published chain is selected as a starting point for auditing. Transactions are then traced backwards through the chain until the subject timestamp is reached. The chain should be unbroken to that point, and the subject timestamp should appear in the correct position in the chain between timestamps occurring before and after. Finally, as an additional verification measure, the published chain information for the period of the timestamp can be obtained (e.g. from newspaper archives), and the cumulative hash for the transactions in that time period can be recomputed and verified to match the published value. This provides further assurance since tampering with newspaper archives is unlikely to be practical.

These same methods, while developed in the 1990s for the simpler purpose of time stamping, form the basis of modern blockchain systems.

2.2.3.2 Structure of a Distributed Ledger

In order to establish a chain within the distributed ledger, the ledger is structured into a format which readily facilitates cryptographic verification. Chaining every transaction on top of the previous transaction would create several problems. First, chaining enforces ordering, but with a large transaction volume in a distributed network ordering is both difficult to determine (due to clock skew between nodes, propagation through the network, etc) and not particularly important (the order of transactions need not be verifiable to the sub-second level as long as the same coins are not spent to closely together).

Instead, transactions are collected over a period of time and formed into a block. Each block is handled as a single unit for the purpose of chaining. This allows a lower number of total operations

to verify the chain, and the longer time interval between each block allows sufficient time for propagation throughout the network.

For individual transactions to be verifiable, it must be possible to efficiently verify that a transaction exists in a block unmodified. This can be done by producing a cumulative cryptographic hash of all transactions in the block, and including that as a portion of the block data which is covered by the chaining scheme. While there are several possible methods of producing this cumulative hash, the method most often used in cryptocurrency is the Merkle tree. Introduced in (?), Merkle trees are a (typically binary) tree structure in which each node contains the hash of its children. These children include the hash of their children, and so on, until the leaves are the data which is to be signed.

The Merkle tree is a simple construction and has a useful property: given the full representation of a transaction, it is possible to confirm that the transaction exists in the block using only the Merkle tree itself (e.g. the values of the nodes), without the other transactions in the block. While currently often unused, this permits a significant optimization in the amount of data required to verify that a transaction exists in a block.

Generally, each block consists of a header which contains metadata and then a Merkle tree, in which the leaves are only references to the full transactions. The rest of the block then consists of those transactions, typically located within the block by offsets.




Figure 2.2 Structure of example cryptocurrency blockchain. Each block refers to the hash (actually Merkle tree root in the Bitcoin implementation) of the previous block.

2.2.4 Proof of Work

2.2.4.1 Selecting Valid Blocks

The blockchain is made up of a series of blocks, each of which contains the hash of the previous block. By the same principle as linked timestamps, the blockchain can be verified, and due to the internal structure of the blocks, individual transactions within each block can be verified. This creates an auditable distributed ledger, so long as the community of users of the cryptocurrency are in agreement about the most recent legitimate block.

This presents a substantial challenge. There are two problems which must be solved: More obviously is the distribution of the legitimate most recent block to the network. However, a number of peer-to-peer technologies capable of solving this problem are well known, included distributed

hash tables and other peer-to-peer state synchronization protocols such as BitTorrent. Second, though, and more difficult, is the selection of a block to serve as the "head" of the blockchain, thus establishing the legitimate chain. Surety solved this problem for their time stamping system by acting as a central authority and selecting the group they knew, by merit of having generated it, to be legitimate. However, this centralized system requires trust. Bitcoin and other cryptocurrencies seek to solve this problem in a completely distributed fashion.

This problem is not only evident in the case of fraud. Blockchain systems require that some node act to gather the most recent transactions into a new block, and then that other nodes accept that block. However, the simple logistics of distributed systems mean that this process will sometimes go wrong. Occasionally, two legitimately acting nodes will each generate a block at the same time. Because the transactions themselves are spread through the distributed system, one of these nodes may be aware of transactions that the other has not yet received. This results in two valid, but different, blocks.

If every node in the network was guaranteed to know about every block which had been generated, they could each select a block as "most correct" using some combination of useful criteria (e.g. contains the largest set of transactions) and criteria intended only to encourage a sole selection (e.g. timestamp with a certain value in its last several bits). The node could then base any future work off of that "most correct" block, which should result in the "less correct" blocks forming a branch which will die off as most or all nodes choose to ignore it.

Of course, at this stage the contents of each block will also be validated for the legitimacy of its transactions. For example, each transaction will be checked on the basis of its cryptographic signatures to ensure that it was issued by someone authorized to spend the coins it consumes. Any block which fails these cryptographic validity checks will be ignored by any legitimate node in the network.

This introduces an important concept in blockchain systems: while the blockchain is nominally a chain, in practice it is a *tree*, with various branches created whenever, for reasons legitimate or fraudulent, two blocks are created chained off of the same previous block. Whenever this occurs, the nodes making up the network must agree on some selection criteria for which block will be preferred. If the majority of nodes in the network agree on which block to select, then even if some nodes had not yet received that block (and thus had started work off of the less preferred block) they will later detect that the remainder of the network has selected a different branch and redirect their attention.

This can simply be referred to as the "longest chain" rule: to manage branching of the blockchain, nodes apply some agreed-upon selection criteria to determine which block to work from. If any branch has had more blocks added to it though, nodes prefer the longest branch, which ensures that "majority rules" when evaluating these criteria. This will be made more clear by an example shortly.

2.2.4.2 Byzantine Generals

There is a problem with this model: it assumes that the nodes in the network each act legitimately, and that no node e.g. generates new blocks at a high rate to artificially produce a new longest chain which includes fraudulent content.

This is closely related to what is referred to as the *Byzantine Generals problem*. Prominently described in (?), this well-known problem in distributed computing relates to the challenges of reliably coming to an agreement when some actors in the system may intend to tamper with the process. Consider this classical formulation:

We imagine that several divisions of the Byzantine army are camped outside an enemy city, each division commanded by its own general. The generals can communicate with one another only by messenger. After observing the enemy, they must decide upon a common plan of action. However, some of the generals may be traitors, trying to prevent the loyal generals from reaching agreement. The generals must have an algorithm to guarantee that

(A) All loyal generals decide upon the same plan of action.

The loyal generals will all do what the algorithm says they should, but the traitors may do anything they wish. The algorithm must guarantee condition A regardless of what the traitors do.

The loyal generals should not only reach agreement, but should agree upon a reasonable plan. We therefore also want to insure that

(B) A small number of traitors cannot cause the loyal generals to adopt a bad plan.

[?)]

An important element of the problem is this: the content of the generals' plan is not particularly important, so long as they agree on it so that they form a unified force. The case in cryptocurrency is very similar. Because transactions are expected to be independently verifiable as legitimate, the concern at the stage of block formation is whether or not a transaction is included. A transaction that was intended but fails to appear in the block can be reattempted for the next block. A transaction that was not intended cannot appear in a block since the verification of its cryptographic signature would fail. Much like the Byzantine generals, the goal of a cryptocurrency network is to agree on a *consistent* blockchain.

?) introduced a novel solution to this problem: *proof of work*. If the generals are only able to send each other messages slowly (more slowly than the time it takes to pass a message to all other generals), then they will have to agree on whichever message can be generated first, since no other messages will be received that conflict with that one.

In the real world, it is difficult to see how to enforce that messages are sent slowly. There is a relatively common solution in computing, however: force each general to solve a computationally difficult problem before sending a message, and to include in the message proof that they completed the computation. This is the approach taken by cryptocurrencies.

We can understand the proof of work function by taking the example of that used by Bitcoin, hashcash. Hashcash was originally introduced in ?) for the purpose of deterring unsolicited email by requiring email senders to prove that significant computation had been exerted to generate the email (thus making it expensive to send large volumes of email). Hashcash is fairly simple. Given a cryptographic one-way function (such as, in the case of Bitcoin, SHA256 performed twice sequentially), find an input to the hash function that yields a result with a specific prefix (such as, in the case of bitcoin, a series of zeroes).

In practice, solving this problem requires a brute-force effort of attempting a large number of inputs until a satisfactory one is found. While the time this task requires is stochastic, the difficulty (such as the length of the required prefix) can be adjusted to set an average amount of computational effort required.

The blockchain system thus depends on both selection of the longest chain, and proof of work to ensure that candidate blocks can be spread through the network quickly enough for the majority of nodes to evaluate them. Consider an example of this process:

Following an accepted block of the blockchain, block A, three nodes each complete the proof of work function, relatively closely together in time, to create a new block. Of these three blocks B, C, and D, B was created by a dishonest node that has attempted to include an invalid transaction. C and D were created by honest nodes, but vary slightly due to differences in transaction propagation time. This results in a fork in the blockchain.

One node on the network, node 1, selects C as the best. Node 1 begins the proof of work computation to produce a new block based on C. At the same time, node 2 selects D and begins the proof of work computation to generate a new block from block D. Node 1 completes the proof of work function more quickly and publishes a new block, E, based on block C. When node 2 receives block E, it validates it and then, seeing that block E has become the longest branch, abandons its efforts on block D and begins to compute a new block based on block E.

Because forks in the blockchain sometimes occur and persist for a short time, cryptocurrency implementations generally do not consider a transaction as final until the blockchain has been extended past it by several blocks. For example, in the case of Bitcoin, it is common to wait for 10 additional blocks, a process that is often described as waiting for ten verifications of the transaction, since each additional block indicates the network accepting that another node has verified the transaction.




Figure 2.3 Structure of the blockchain, showing a fork which has been abandoned by miners after another fork overtook it in length, ensuring a single longest chain.

2.2.5 Compensating Computation

This entire process is computationally expensive. First, the process of validating a block for correctness is intensive since it requires performing a large number of cryptographic operations. Second, the proof of work function is time-consuming. Indeed, a full participant in a cryptocurrency network will generally be continuously engaged in proof of work computation, since it always attempts to create a new block based on the longest chain. Ultimately, the majority of that computation is wasted, since it is unlikely that any given node will be the first to complete the proof of work.

This computational effort amounts to very real cost in the form of electricity and cooling, even when done at a relatively small scale. As a result, the system is designed to compensate the nodes which perform this computation.

There are two methods of compensation: transaction fees and mining fees. First, each transaction typically includes a small amount of currency “input” without a corresponding “output”. This difference is the transaction fee, and is included as an incentive for nodes generating new blocks to include that transaction in the block. Whichever node successfully generates a block including that transaction is entitled to keep the total of all transaction fees in the block. As a result, transactions

with low or no fees may be ignored by nodes which do not find it worth it to expend the effort on validating them.

Mining fees are intrinsic to each block: any node that successfully generates a block is entitled to generate anew a certain amount of currency as the block mining fee. This term “mining” is used in reference to gold mining, and the process of verifying transactions and attempting the proof of work to generate a new block is broadly referred to as cryptocurrency mining, and the nodes that do it as miners. Cryptocurrencies generally do not require a node to mine, but doing so is potentially profitable for the node’s owner.

In addition to incentivizing verification of transactions, the block mining fee is the mechanism by which *new* coins are introduced. In most, but not all, cryptocurrencies, the value of the block mining reward decreases over time. This is a feature intended to make cryptocurrencies inherently deflationary (that is, increasing in value), rather than possibly inflationary as would be the case if new coins were added at the same rate indefinitely.

2.3 Bitcoin

2.3.1 Bitcoin Standards

The specific implementation of Bitcoin will now be discussed. It is worth prefacing this material with the history of Bitcoin’s nature as a technical standard, which has been peculiar and sometimes controversial.

Bitcoin’s seminal publication, (?), describes the system at a high level but provides minimal technical detail. It was released pseudonymously, and besides some interaction with the nascent Bitcoin community early on the author has remained silent. Many attempts by the press and the community to identify the author have been unsuccessful [see, e.g., ?]. While this adds greatly to the intrigue surrounding Bitcoin, it has also led to challenges in the standardization and maintenance of Bitcoin due to the lack of clear leadership early on in the system’s evolution.

In addition to the paper, Nakamoto released an implementation of the described system, including source, written in C++. Due to the lack of a detailed written specification, this implementation became the de facto standard, and in order to maintain compatibility all subsequent implementations have copied all aspects of the inter-node communication protocol and the blockchain directly from that implementation. Additionally, the most common implementation of Bitcoin today, Bitcoin Core, is descended from the Nakamoto code base.

There are several eccentricities of the Nakamoto code base that are now encoded into the protocol and could serve as useful signatures of Nakamoto’s authorship. For example, the cryptographic hashing function used throughout is SHA256, but curiously, the algorithm is always run twice in sequence, and the bytes of the resulting hash are displayed and stored on disk in the opposite of the conventional order. This convention, while irritating to developers working to maintain compatibility with existing Bitcoin implementations, has remained unchanged and is generally referred to as SHA256d.

As a second example, the Bitcoin implementation makes extensive use of an unusual variable-length integer representation referred to as a *CompactSize*. However, this format is appreciably less efficient than other well-known variable length integer encodings such as Base128. The Bitcoin Core implementation has switched from CompactSize to Base128 in data structures not communicated to other nodes, but to preserve compatibility CompactInt must still be used over the network and in data structures subject to hashes or signatures.

Changes to the Bitcoin protocol are rare because they require adopters of the new version to form a new network independent of users of the old protocol version, which will ignore communications in the new version as invalid. Because the transition between versions cannot occur instantaneously, this leads to an event referred to as a *hard fork* in which the Bitcoin network fractures into two separate networks using the two separate versions. Because all currency that existed before the fork will exist in both versions of the network after, this creates a potentially chaotic situation in which Bitcoin in both the old and new network may still have monetary value. For this reason, hard forks are generally viewed as completely new cryptocurrency systems and operate under a different name. Bitcoin Cash, Bitcoin Gold, and Bitcoin SV are all examples of currencies hard-forked from Bitcoin². Each of these forks were both technically and, perhaps more importantly, politically complex [see, e.g., ?].

Further complicating matters, there have been two unintentional hard forks of Bitcoin resulting from a protocol change [?] and one a security vulnerability [?] in the design of the protocol. In both cases the pre-update version of Bitcoin has fallen out of usage, and in one case it was intentionally resolved by use of a later soft fork, but these incidents have raised concerns about the stability of the Bitcoin protocol.

Finally, Bitcoin has gone through a number of "soft forks." In a soft fork, blocks and transactions using the new protocol are still recognized as valid by clients implementing the older protocol (although older clients may not be capable of generating them). This prevents a duplicate situation, since the entire network will remain in agreement about valid blocks. Soft forks have happened intentionally a number of times and have generally gone smoothly, although there has been one incident (that of BIP66) in which a soft fork resulted in a brief near-hard-fork condition due to implementation inconsistencies [?].

All of this is to say that precise standardization of the Bitcoin protocol is based on consensus of the blockchain (e.g., is derived from the contents of the current longest chain), rather than from written documentation prepared by a conventional standards process. While the community has worked to introduce a high level of standardization through the Bitcoin Improvement Proposal (BIP) process, changes in the Bitcoin system are often as political as they are technical—if not more so. Many apparent oddities of the implementation today are a result of this process.

2.3.2 Bitcoin Transactions

Throughout discussion of the Bitcoin implementation, detailed protocol descriptions (e.g. byte counts and byte offsets of fields) will be omitted. The Bitcoin protocol is documented in such detail in the Bitcoin Wiki³ which is generally considered an accurate reference, although not a prescriptive standards document. Instead, enough information will be given to illustrate how the details of the Bitcoin protocol and implementation impact forensic analysis.

Each Bitcoin transaction is preceded with a version number. Four bytes are reserved for this version number, although it has never been incremented beyond 1, introducing the possibility for reuse of these fields depending on client handling of unrecognized versions. A flag field indicates the use of SegWit (discussed later), followed by a list of inputs, a list of outputs, and a list of

²The fact that two of these soft forks originated over disagreement on the target size of each Bitcoin block—1MB with segmented witness (Bitcoin), 32MB (Bitcoin Cash), or 128MB (Bitcoin Gold)—shows how difficult technical decisions can become in the decentralized currency context. Bitcoin Gold originated from a change in the proof of work protocol to resist ASIC acceleration, a matter that has led to a large number of altcoins, although others started from a new genesis or first block rather than hard forking from Bitcoin.

³<https://en.bitcoin.it/wiki>

witnesses if SegWit is in use. Finally, each transaction contains a "lock time" field which prohibits its being recorded in the blockchain until a certain amount of time elapses. This field is mostly used in conjunction with scripts, discussed later.

Each input in the input list consists of a reference to the previous transaction which it is spending (by the hash of that transaction), the index of the output within that transaction that is being spent (different outputs of a transaction are often sent to different people, and so most likely only one specific output is eligible to be spent), an input script, and finally a sequence number used as part of the lock time feature.

Outputs are somewhat simpler, consisting of a value field giving the amount of the output and an output script.

Bitcoin scripts are short programs written in a constrained stack-based language with semantics similar to Forth. Bitcoin scripts do not provide loops, avoiding nodes verifying transactions becoming stuck on a script which does not complete execution. Available script operators include simple logic (e.g. checking equality) and cryptographic operations such as hashing⁴ and signature verification. Each operation is specified as a one-byte opcode, while opcodes are provided to easily handle data values from 1 to 75 bytes.

It may be confusing that inputs and outputs consist only of a script, particularly since Bitcoin is typically discussed as transferring funds between keypairs. In fact, these scripts are the entire basis of transaction verification in Bitcoin. Verification of a transaction is performed by first executing the *input* script on the spending transaction, and then executing the *output* script of the funding transaction with the same stack.

The flexibility that this creates is one of the most powerful features of the system. It is easier to understand, though, if we start by discussing the simplest transaction type: payment to a public key. Often referred to as a P2PKH transaction (Pay to Public Key Hash), this simple and extremely common type of transaction leads to the output script often being referred to as the "ScriptPubKey" and the input script as the "ScriptSig," in reference to the role they serve in a P2PKH transaction. However, it should be noted that these names are only conventional due to a common use of the scripts and do not dictate how the scripts should be used.

When a user intends to simply send Bitcoin to another user, they generate a P2PKH transaction. This transaction has an output directed at the intended recipient, with a script of the form `OP_DUP OP_HASH160 (hash of recipient public key) OP_EQUALVERIFY OP_CHECKSIG`. The input script to spend this transaction takes the form *(signature) (public key)*.

When executed, the two values of the input script are simply pushed onto the stack. The first operation of the output script then duplicates the bottom item of the stack, the public key, followed by hashing it. The public key recipient hash is then pushed onto the stack again by the output, and the two are verified to be equal. This checks that the public key of the recipient is the same public key intended by the sender. Finally, with just the original signature and public key from the input remaining on the stack, the signature is checked for validity. The `OP_CHECKSIG` operation verifies that a signature *over the entire transaction* is valid against the public key below it on the stack. This ensures that the recipient public key has signed the transaction.

This can be visualized in a table matching the format used in the Bitcoin Wiki, Table 2.1. This format shows the stack contents and remaining instructions (scripts are executed in strict left to right order) after each step of evaluation.

⁴It is a further reflection of Nakamoto's eccentric approach to cryptographic algorithms that two hashing functions are available: `OP_HASH160` which performs SHA256 followed by RIPEMD160, and `OP_HASH256`, which performs SHA256 followed by SHA256. The latter is rarely used.

Along with the simple P2PKH transaction, scripts also allow a number of more complicated transaction types, such as transactions that require multiple signatures from different users. These “multisig” transactions allow a single transaction to have inputs from multiple users. Since all transactions can have outputs to multiple users, multisig transactions serve as many-to-many exchanges which can make the flow of currency exceptionally difficult to track.

One of the most complicated transaction types is a relatively new feature of the system: pay-to-script-hash or P2SH. In a P2SH transaction, the output provides only the script `OP_HASH160 hash of script OP_EQUAL`. This appears to be just a verification that the input script is able to produce some arbitrary value. In fact, P2SH transactions, by merit of having this output script, trigger special handling within the Bitcoin implementation that changes the conventional ordering and runs the input script, followed by the output script, followed by the input script again. The input script is expected to leave one value on the top of the stack, which is itself a script which is then executed.

This rather eccentric parsing process enables the most interesting feature of P2SH transactions: A P2SH output specifies the hash of a script that will satisfy it, but it is up to the *input* to provide that script. The address of a P2SH transaction is actually the hash of a script, not a user. This enables a number of interesting features, such as improved multisignature support and the ability for a transaction to be sent to a script without knowledge of the content of that script, only its hash. P2SH transactions are becoming increasingly common for a variety of applications now ranging from simple multisignature control to smart contracts.

Scripts can vary greatly in length, particularly if they include multiple signatures, which are relatively long. As a result, Bitcoin transactions themselves vary in size. Because larger transactions take longer to verify (due to script execution time) and take up more space in blocks, the system is designed to provide an incentive to keep transactions smaller. The fee on a transaction is calculated on a per-byte basis. As a result, larger transactions must include a larger fee in order to assure inclusion by miners.

This aspect of Bitcoin—both the fees on larger transactions and their greater impact on blocks—has led to a great deal of pressure to modify Bitcoin to enable high volumes of more complex transactions. This pressure has been a key flashpoint resulting in forks of Bitcoin and the development of altcoins. In the case of mainline Bitcoin (that is, not Bitcoin Cash or Bitcoin SV), the Segmented Witness or SegWit solution has been implemented. This is a modification to the block structure.

2.3.3 Bitcoin Blocks

Bitcoin blocks consist of a header followed by a list of all transactions in the block. The header contains the hash of the previous block, the hash of the root node of a Merkle tree representing the block, a timestamp, and a “target” value which sets the difficulty of the proof of work function.

By agreement of the miner implementation, the target is dynamically adjusted so that a new block is produced, on average, every ten minutes. This adjustment is referred to as a “retarget” and the calculation is made by all miners approximately every two weeks (established by counting the number of elapsed blocks since the last retarget). Blocks with improperly calculated targets will be ignored as invalid.

The target value itself is calculated based on a compact floating-point value in the block header referred to as the bits or nBits. A proof of work computation is considered successful if it yields a hash which is less than the target. The value used for the proof of work is actually the merkle root itself, and so an additional nonce value is stored in the header

The Merkle tree is produced using a simple algorithm which is deterministic given the order of the transactions within the block. For this reason, the block is stored without the Merkle tree, only with the transactions and the Merkle tree root value for faster lookup. It is possible, although rarely implemented, for a node possessing the full block to calculate the entire Merkle tree and send only the tree to another node, allowing that second "partial" node to validate the presence of single transactions in the block without the need for the entire block data.

To compose the merkle tree, the entire set of transactions is taken in order to form the bottom-most row. The next row is composed of the hash of the hashes of the two transactions under it, halving the size of the row. At each step, if there is an odd number of items in the row, the last item is duplicated to ensure a binary tree. This process is repeated until a single item is reached, which is the Merkle tree root. As usual, the hash function used is doubled SHA256.

The scripts included on inputs (ScriptHash) are sometimes referred to as the *witness data* since they are required to confirm a transaction as valid, but not to understand the result of the transaction (which can be inferred based only on the input references to previous transactions). SegWit was a soft fork which modified the protocol to change the storage format of transactions. SegWit-enabled blocks, which are now the only acceptable blocks, omit the witness data from the set of transactions and instead store them separately at the end of the block.

Additionally, the formula used to calculate the size of a transaction (for purposes of establishing the fee required) was modified so that the segregated witness data is counted at one fourth the rate of the transaction itself. More specifically, the unit of transaction size was changed to one quarter byte, and the transaction is counted as four units per byte while the SegWit is counted as one unit per byte.

These changes were made to permit larger transactions with more complex scripts without significantly increasing required transaction fees or greatly limiting the total possible volume of transactions, since blocks have a limited size and are created at a limited rate⁵.

As a result, SegWit-enabled blocks (indicated by a flag in the block header) have a body consisting of a series of transactions, followed by a series of script blocks.

The Bitcoin Core implementation stores the blocks in block files. Each block file contains a series of blocks in the order in which they were received from the network. Once the file reaches 128 MB in size, it is closed and a new file is begun. The result is a large folder of block files, with each block being contained somewhere in one of those block files. It is up to Bitcoin Core to skip through block headers in order to maintain the current longest chain in memory.

2.3.4 Reference Databases

In order for a Bitcoin client to be useful, it must keep track of several data structures that are not actually part of the block: the UTXO and, in the case of miners, the transaction index.

The unspent transaction pool, widely referred to as UTXO, is a list of all transactions with outputs that have so far not been consumed by inputs. Bitcoin never explicitly stores the balance of any given Bitcoin wallet, instead, the balance of a bitcoin wallet is the sum of the UTXO associated with that address. As a result, bitcoin clients (including non-mining clients often referred to as

⁵Another obvious approach to this problem is to increase the block size limit, which is the approach taken by hard forks Bitcoin Cash and Bitcoin SV. The advantages and disadvantages of those changes are largely a matter of philosophy and politics, although the SegWit approach has the minor advantage that it is possible for nodes not interested in cryptographically verifying transactions to parse through the entire set of transactions in a block more quickly

wallets) must keep track of the UTXO in order to calculate the user’s available balance and select available outputs to use as inputs for new transactions.

The UTXO is calculated as blocks are parsed and received, and is stored in a flat file database format called LevelDB. LevelDB, originally developed by [Google](#), is a simple key-value store that offers high performance for the storage and lookup of short byte arrays. It has the additional feature of intrinsic support for sorting of keys. Neither of these features are significantly used in the calculation of the UTXO, which is instead based mostly on simple iteration through all unspent outputs.

LevelDB is also used for the maintenance of the transaction index. Because the transaction index is costly to generate and physically large, it is usually disabled on clients that are not involved in mining. However, it is necessary to validate a transaction’s legitimacy, so it will be maintained by miners.

The transaction index links a transaction hash to the block file that contains it and its offset within that file. This allows the client to find the output that resulted in a given input in a newer transaction, a necessary step for executing the scripts for verification.

2.4 User Experience

The technical implementation of Bitcoin is surprisingly complex, and so bears some separation from the experience of using Bitcoin for most users. It is worth explaining a common user experience with Bitcoin, while noting that there are many variations in the tools and services available. For example, this example will omit hybrid wallets, bitcoin exchanges acting as wallet services, and off-chain transactions for simplicity.

A user begins their interaction with Bitcoin by installing a Bitcoin wallet, typically Bitcoin Core, which spends some time (often several days) synchronizing the entire blockchain. About two hundred gigabytes of storage are currently required to store the entire blockchain, and ongoing internet access is required to keep the blockchain state synchronized. These requirements mean that users of mobile devices and other low-power devices will usually not use a full Bitcoin wallet but instead some kind of wallet service which synchronizes the blockchain state externally.

The user’s wallet will then generate a key pair to be used for verifying and creating transactions. This key pair is shown in a short base-58 encoded format referred to as an address. By giving another person or service that address, they can receive coins. An incoming transaction will usually appear in the user’s wallet after a single block confirmation, within ten minutes, although wallets usually show the incoming transaction as tentative or unverified until some number (often six) of additional blocks have been generated “on top of” the transaction. Internally, this incoming transaction is now part of the UTXO belonging to the user, and so its value is added to the display balance.

To send Bitcoin, the user enters a destination address and an amount. The user need not concern themselves with which specific transactions will be used as inputs or with the change output as these are generated automatically (although some wallets do allow the user to select these parameters manually if they so choose). Once they save the transaction, it appears to them as an unverified outgoing transaction until it has been confirmed in a block.

A large portion of users may interact with Bitcoin using only these simple transactions. Because of their technical complexity, more involved transaction types like multisignature or P2SH transactions are generally created with the help of some type of service, such as a website. One of the major advantages of P2SH transactions is that they allow such a service to generate a complex transaction and then have individuals send funds to it from their own wallets. This largely hides the complexity of such transactions from the user.

Table 2.1 Evaluating a P2PKH Script

Stack	Script Operations
	(sig) (pk)
Loaded script from spending input (ScriptSig)	
Loaded script from output being spent (ScriptPubKey)	(sig) (pk) OP_DUP OP_HASH160 (pkh1) OP_EQUALVERIFY OP_CHECKSIG
Pushed literal values on to stack	(sig) (pk) OP_DUP OP_HASH160 (pkh1) OP_EQUALVERIFY OP_CHECKSIG
Duplicated value on top of stack	(sig) (pk) (pk2) OP_HASH160 (pkh1) OP_EQUALVERIFY OP_CHECKSIG
Computed hash of value on top of stack and replaced, yielding expected spender public key	(sig) (pk) (pkhash2) (pkhash1) OP_EQUALVERIFY OP_CHECKSIG
Pushed literal values on to stack	(sig) (pk) (pkhash2) (pkhash1) OP_EQUALVERIFY OP_CHECKSIG
Checked top two items of stack for equality and removed, verifying that signature key matches expected spender	(sig) (pk) OP_CHECKSIG
Verified that transaction signature was issued by spender's public key	

APPENDIX 3. ANTI MONEY LAUNDERING

3.1 Motivations and Methods

Money laundering, broadly defined, is any attempt to conceal the origin of funds that a person has access to. In practice, money laundering is generally conducted for one of two reasons: In some cases, a person has access to funds which they have obtained illegally, and so they need to prevent their association with illegal activity being proven by the flow of money. In other cases, a person obtains money illegally or legally but wishes to conceal their possession of the money, for example to avoid tax liability. Money laundering is also conducted in some cases to avoid tariffs or other types of regulation.

Money laundering has a long history, although it is generally considered to have become prominent around the time of prohibition in the United States [?]. Prohibition resulted in the abrupt formation of a strong system of organized crime which operated quite profitably. The beneficiaries of illegal alcohol distribution naturally wished to take advantage of their wealth, but for law enforcement unexplained wealth on the part of a person associated with alcohol production or sale was an obvious indication of illegal activity.

For a person benefiting financially from crime then, it was best to find or develop a “cover story” which created the appearance that one’s wealth had originated from legitimate business. This remains a common element of organized crime today, particularly with the enhanced ability of law enforcement investigators to conduct statistical analysis to reveal unusual spending.

The motivations and methods of money laundering are explored more formally in (?). However, it is particularly useful in our context to examine some common methods by which money is laundered.

Cash has the significant advantage of anonymity: typically, it is not possible to trace the origin of cash money, which breaks the chain of possession of funds and allows for some degree of money laundering in its own right. Indeed, the movement of huge quantities of cash remains a major part of organized criminal activity [?].

However, in practice simply transferring large amounts of money in the form of cash is insufficient for purposes of laundering, because acquiring a large amount of cash is itself suspicious and can result in heightened attention. Ideally, the cash must be made to appear to be the proceeds of some legitimate enterprise.

This leads naturally to one of the most obvious and common forms of money laundering: injection into a cash-heavy business. A person wishing to launder money can purchase a business which operates primarily in cash, such as a convenience store or nightclub. Proceeds of criminal activity are then injected into that business as if they were normal revenue, as simply as by adding the cash to tills or drop safes before close of business counting.

This type of money laundering, famously employed using a car wash in the television series *Breaking Bad*, is extremely common and effective due to the difficulty of auditing any significant portion of cash-based businesses. To detect such laundering, an investigator would have to observe the operations of a business in order to estimate their actual revenue and then show a substantial gap between their reported revenue and the greatest reasonable revenue considering the volume

of customers and sales. Even a very large difference may be difficult to conclusively prove as illegitimate in examples such as nightclubs where cash purchases are made at a variety of different points in the business and under conditions which make observation difficult.

Further, money launderers often diversify and rotate their so-called fronts, buying and selling cash-heavy businesses throughout an area in order to prevent thorough analysis of any one front during the time period it is in operation. This strategy tends to keep launderers “one step ahead” of efforts to uncover the origin of their money.

Another common method of money laundering is to substitute cash for some type of asset which is easy to purchase and sell. Common examples include precious metals and lottery tickets. The use of lottery tickets may, at first, seem counter-intuitive, since the lottery results on average in a substantial loss. However, lottery tickets are generally extremely easy to redeem with few questions asked, and the large quantity of tickets purchased in order to find large wins simply goes unnoticed. These advantages make even the loss of large-scale gambling reasonable as a money laundering method.

Precious metals are less popular as a means of money laundering because, like cash, the possession of large quantities of gold tends to raise suspicions. However, excuses can often be formulated such as family savings or speculative investment in gold.

From these examples, we can identify two general requirements for the laundering of money: First, it must be impossible to trace the origin of money, for example because it is exchanged in the form of physical tokens that are not generally tracked (this includes cash and precious metals). Second, the money resulting from the laundering process should appear at least possibly legitimately acquired—enough to establish plausible deniability (for example, operation of a business or lottery winnings).

As discussed in [?], it is very difficult to determine how commonly money laundering is performed or at what scale. It is, after all, intended to be an undetectable, untraceable activity. However, some estimates are attempted in [?]. For the twenty OECD countries, US\$603 billion may have been laundered in the year 2006 alone. Money laundering is also understood to have a significant nexus with organized crime, laundering serving as a key mechanism which allows large-scale criminal operations to sustain themselves.

Money laundering was fast to move online. [?] discusses conventional methods of laundering money online, which include the extensive use of online casinos and betting services and smaller-scale abuses of money transfer services such as PayPal, particularly when combined with prepaid debit cards.

[?] also mentions an interesting example of prior art in digital currency: liberty reserve (LR). LR was a centralized digital currency service similar in concept to PayPal, but deposits were “exchanged” to a currency called LR Dollars or LR Euros, depending on whether the original deposit was made in US Dollars or Euros. LR did not require any identification from its users and maintained no anti-money-laundering program, which made it particularly attractive to criminal users.

In 2013, US law enforcement arrested the operators of Liberty Reserve and seized the website. This action had been taken in response to significant criminal usage of the service and questions over its funding and possible involvement in fraud. In [?], security industry reporter Brian Krebs suggests that the seizure of the Liberty Reserve service caused a significant disruption to organized internet crime.

3.2 Money Laundering and Computer Crime

Computer methods of money laundering have a natural nexus with computer crime, since they allow profitable criminal ventures to operate in a more automated, lower-friction environment. An especially clear example is that of “ransomware.”

Ransomware is a family of malware including such notable instances as Cryptolocker and Wannacry. These programs encrypt (or at least obfuscate) many of the files on a victim’s computer and then demand a ransom in order to decrypt them. For the many individuals and businesses that lack adequate backups, this can lead to a devastating loss of records and information that makes victims likely to pay the ransom.

According to [?], more than 40 US municipal governments have been struck by ransomware. The resulting damage is so extensive and costly that one municipal government, that of Lake City, Florida, paid a ransom of nearly a half million dollars in order to regain access to their data. Similar incidents occur in hospitals, school systems, corporations, and the homes of individual consumers, which may be all the more likely to pay the ransom and leave the matter entirely unreported.

This profitable scheme presents a challenge, though: how to receive the ransom funds without making oneself subject to prosecution. This is particularly important since bold attacks on governments and institutions have attracted the attention of numerous law enforcement agencies across countries and levels of government.

Ransomware has a surprisingly long history. However, early examples of ransomware were hampered by relatively complex and risky payment schemes which ranged from premium-rate telephone numbers to overseas online pharmacies serving as mules. [?] suggests that the increasing prevalence and damage caused by ransomware is due in large part to the availability of easier to use, more reliable, and more anonymous payment channels.

[?] investigated the historical development of ransomware starting in 1989. A notable finding is that since approximately the introduction of Cryptolocker in 2013 (often seen as the first massively successful ransomware), the use of cryptocurrency as the payment mechanism has been a virtual requirement for ransomware.

Because of the nature of the distributed ledger, however, ransom payments made via cryptocurrencies can likely be traced by law enforcement if they are not processed through some type of laundering service. These laundering services are a key enabler of ransomware and many other types of computer crime.

3.3 The Bank Secrecy Act

3.3.1 Reporting

US efforts to counteract money laundering gained their present shape in 1970 with the passage of the Bank Secrecy Act (BSA). The BSA was the first step in transferring the prevention of money laundering from an entirely post-hoc investigative activity conducted by law enforcement to a regulatory scheme with controls enforced by financial institutions.

The BSA faced legal challenges, most notably in *California Bankers Assn. v. Shultz*, 416 US 21, before the Supreme Court in 1974. These argued that the BSA constituted a violation of either due process or search and seizure protections, since it required financial institutions to actively monitor customer activities before the beginning of any criminal process. However, the BSA withstood these challenges and remains in place as the core of AML regulations today.

The key provisions of the BSA are requirements that financial institutions file several different types of reports on their customer's activity. These reports include the well-known currency transaction report, or CTR, which is a report of any movement of more than \$10,000. Whenever a customer transacts more than that amount, even if across multiple transactions or otherwise obfuscated, the financial institution must report that activity to the government.

Financial institutions must also maintain records on the purchase of any monetary instrument valued greater than \$3,000 (but less than \$10,000, since greater amounts would result in a CTR) in what is known as the monetary instrument log or MIL. Although not submitted to the government, these records are retained for five years for inspection by law enforcement.

Perhaps most interestingly, the BSA also mandates suspicious activity reports or SARs. SARs must be filed whenever a customer appears to be acting suspiciously. The most clear case of suspicious behavior which results in a SAR is "structuring," the use of multiple transactions to avoid activity becoming subject to CTR or MIL. For example, depositing \$5,000 each on three subsequent days to avoid reaching the \$10,000 threshold for a CTR is a simple case of structuring which will result in a SAR. More complex cases are also possible, though, and financial institutions are more broadly required to report any suspicious activity they are aware of.

3.3.2 Due Diligence

Further, financial institutions are required to actively identify suspicious activity. Amendments to the BSA, notably the Money Laundering Suppression Act of 1994 and the Money Laundering and Financial Crimes Strategy Act of 1998, require financial institutions to develop a comprehensive AML program that includes training of employees to recognize suspicious activity and to establish AML examining, essentially active auditing of customers for potential money laundering. AML examination may be conducted by a dedicated department in larger institutions or as one of several duties of bankers and auditors in smaller institutions.

Title III of the Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act (USA PATRIOT) of 2001 introduced the most significant enhancement in US AML regulation since the BSA. USA PATRIOT requires that banks conduct due diligence on their customers to ensure that they are not involved in criminal activity. This due diligence requirement forms the most significant part of the broader framework of "know your customer" or KYC, discussed later.

3.3.3 Anomaly Detection

Financial institutions can face harsh penalties for failure to report suspicious activity of which they should have been aware. Both civil and criminal sanctions are possible, and civil fines have run as high as \$160 million [?]. As a result, financial institutions are motivated to detect suspicious activity to prevent possible penalties.

Financial institutions employ a variety of methods to detect suspicious activity and comply with reporting requirements. As with most aspects of the financial industry, this responsibility is highly automated. The industry association report (?) discusses best practices in automation of AML compliance.

Most automation is computationally simple and includes basic activities like maintaining records of customer identification and querying for transactions which match a list of terms associated with illegal activity (i.e. in their recipient names).

Other methods of automated analysis are more sophisticated. Statistical analysis of transactions to detect outliers and changing trends is common, and many such systems now develop per-

customer profiles for more sensitive alerting on changes. These systems may fuse information from not only a bank's own records but also from business partners and intelligence sharing communities, providing a complex data set to evaluate.

As a result, complex statistical approaches to money-laundering detection are familiar to large financial institutions, which have typically already devoted substantial resources to development in this area.

3.3.4 Scope

The BSA is larger in scope than may be obvious. BSA was drafted with an eye towards depository financial institutions such as banks primarily, but is intended to cover many types of organizations which would be a component of a money laundering process.

First, BSA reporting requirements generally apply to the individual transacting as well as to financial institutions. When conducting a transaction of greater than \$10,000, it is the responsibility of the individual under the law to ensure that a CTR is submitted to FinCEN, even though they made need to do so themselves. Individuals and organizations are required to report directly to FinCEN any holdings of greater than \$10,000 in foreign financial institutions as well.

The exception to consumer involvement in reporting is the SAR. In the case of SARs, the customer is specifically forbidden from knowing that an institution has submitted a SAR related to their financial activity. Institutions may be penalized for allowing a customer to learn about a SAR related to their actions.

CTR reporting requirements apply not only to financial institutions but to any organization that may be involved in large transactions. Within the financial industry, this includes investment brokerages and financial services companies. Outside of the finance industry it includes casinos, car dealerships, pawn shops, and more.

Such non-financial institutions typically have less complex AML programs, including less analysis on transaction histories. However, they are subject to similar due diligence and KYC requirements, and so may employ many of the same tools and methods.

3.3.5 AML Regulators

US financial institutions operate within a complex regulatory environment with multiple authorities of varying jurisdictions. However, anti-money-laundering requirements rest largely with the Secretary of the Treasury who, since its establishment in 1990, has delegated them to the Financial Crimes Enforcement Network or FinCEN. FinCEN, now a bureau of the Treasury Department, acts as the US "financial intelligence unit" in collecting and analyzing BSA and other reports for unusual activity [?].

A number of other government or semi-governmental bodies have regulatory duties which affect AML enforcement, including the three banking regulators (the Board of Governors of the Federal Reserve System, the Office of the Comptroller of the Currency within the Treasury Department, and the Federal Deposit Insurance Corporation). Other types of financial services institutions may be regulated by the National Credit Union Administration, the Securities and Exchange Commission, the Internal Revenue Service, some stock exchanges with authority to self-regulate, and so on.

Further, many organizations with AML obligations may not have a financial regulator per se at all, but are still bound to issue reports to FinCEN. This situation is typical of retail businesses and others that are not necessarily financial institutions but are involved in transactions reaching the CTR reporting threshold.

3.4 Know your Customer

3.4.1 KYC Principles

The crux of BSA, and AML regulation more broadly, is due diligence requirements often referred to as KYC. Institutions subject to AML regulation are expected to perform KYC on their customers prior to doing business and on an ongoing basis. Further, they are expected to perform KYC on the beneficial owner of customer accounts and monetary instruments, when that person is not the same as the direct customer. With some exceptions the law does not lay out hard requirements for KYC, instead requiring only that it be sufficient and in line with industry norms. Instead, the industry has developed a common set of practices based on experience.

The simplest form of KYC, which is now nearly universal for even smaller financial transactions, is a customer identification program (CIP). A CIP consists of reasonable measures taken to establish the true identity of each customer, such as inspection of a photo ID, and retention of records showing the identity of each customer and how it was established. A CIP is necessary in order to comply with BSA filing requirements, and due to ongoing fraud has become common even for small-dollar services such as prepaid debit cards.

Following the USA PATRIOT act, the Treasury Department promulgated standardized requirements for a CIP. These requirements are relatively loose, but require that depository financial institutions verify the identity of their customers, retain a record of that identity, and ensure that each customer does not appear on a list of persons suspected to be involved in terrorism (the Specially Designated Nationals And Blocked Persons List, or SDN).

Further, a KYC program typically involves researching the customer's identity against a variety of outside data sources. For example, a customer which is a prominent elected official may be referred to as a "politically exposed person" and is assumed to present a higher risk of involvement in bribery and money laundering. Similarly, a person with a criminal background represents a higher risk of financial crime due to the possibility of recidivism.

Further, a customer's known business and financial connections may be considered to determine if they are associated with other individuals representing a high risk. This is particularly significant as the financial services industry continues to consolidate, meaning that financial institutions often have a rich history of previous business with a customer and their close associates (e.g. business partners, family members) which provides more information on their relationships.

These factors are all combined to create a risk profile for a customer. This risk profile is an evaluation of the risk the customer represents to the financial institution of involvement in various types of malfeasance. A customer who poses too significant of a risk may be turned away, similarly to the way that a loan underwriter may decline a loan to a person thought to be likely to default. A customer which poses an elevated risk who is still permitted to do business will generally be subject to an enhanced level of monitoring and stricter criteria for alerting and regulatory reports.

KYC also involves an ongoing monitoring component, in which institutions will continue to research customers and compare their own customer databases against incoming information about known and suspected bad actors.

The USA PATRIOT Act further formalized the practice known as enhanced due diligence or EDD. EDD consists of more careful monitoring of customers who are known to present a higher risk, such as those holding significant assets overseas. EDD has further complicated KYC programs as it both increases the level of scrutiny applied to customers and provides a new course of action for customers thought to undergo a sudden change in behavior.

3.4.2 KYC Risk Factors

Standardization of KYC practices is loose, and most of the details of KYC are left to individual institutions to develop based on their own business practices and risk acceptance. Globally, legal requirements on KYC are highly varied. It is reflective of the complexity and inconsistency of global KYC regulations that one global “quick reference” on KYC requirements spans nearly 700 pages [?].

Review of a customer is most complex when that customer is a corporation or other legal entity rather than a natural person. In this case, institutions are required under the USA PATRIOT act to determine the beneficial owner(s) of that organization. A beneficial owner is defined legally as being *at least* any person with greater than 25% ownership, but institutions are generally expected to use lower thresholds in higher-risk cases.

In evaluating the risk a customer presents, institutions are generally expected to consider the complexity of the customer’s ownership structure, the customer’s home jurisdiction and the quality of its AML measures, the degree to which cash is involved in the customer’s operations, the regulatory environment in which the customer operates, and more [?].

The Federal Financial Institutions Examination Council, a joint council of multiple US financial regulators, states that a KYC program include:

- Obtaining and analyzing sufficient customer information to understand the nature and purpose of customer relationships for the purpose of developing a customer risk profile; and
- Conducting ongoing monitoring to identify and report suspicious transactions and, on a risk basis, to maintain and update customer information, including information regarding the beneficial owner(s) of legal entity customers. [?].i

3.5 AML and Bitcoin

3.5.1 Applying Regulation to Cryptocurrency

It should be clear that AML regulations pose a significant challenge to businesses providing financial services to Bitcoin users and potentially even conducting non-financial business in Bitcoin.

In general, Bitcoin poses a substantial challenge to regulators. It is significant that even the basic legality of Bitcoin in consideration of the Stamp Payments Act is somewhat questionable. ?) discusses the patchwork legal analysis which has been applied to Bitcoin so far in contexts ranging from its basic legality to consumer protection law.

This concern is not theoretical, and while the Bitcoin industry has been slow to adapt, it has been clear for some time that AML regulations posed a significant concern. In ?), FinCEN issued a final ruling establishing that items and values which could be substituted for currency were considered the same as currency for purposes of AML regulation of money services businesses.

In ?), regulatory guidance was issued introducing the term “convertible virtual currency” (CVC) and stating that CVCs are subject to FinCEN regulation:

FinCEN’s regulations define currency (also referred to as “real” currency) as “the coin and paper money of the United States or of any other country that [i] is designated as legal tender and that [ii] circulates and [iii] is customarily used and accepted as a medium of exchange in the country of issuance.” In contrast to real currency, “virtual” currency is a medium of exchange that operates like a currency in some environments, but does not have all the attributes of real currency. In particular, virtual currency does

not have legal tender status in any jurisdiction. This guidance addresses “convertible” virtual currency. This type of virtual currency either has an equivalent value in real currency, or acts as a substitute for real currency.

FinCEN’s position is that virtual currencies are subject to FinCEN regulation because, like instruments of money transmission (i.e. money orders) and other types of issued units of value, they can be converted to and from real currency, making them effectively means of exchanging currency.

The 2013 guidance further states that “The definition of a money transmitter does not differentiate between real currencies and convertible virtual currencies. Accepting and transmitting anything of value that substitutes for currency makes a person a money transmitter under the regulations implementing the BSA.”

Registered money transmitters such as Western Union or the United States Postal Service are in the business of issuing documents which substitute for currency, since they are purchased with currency and then are directly redeemable for currency—consider, for example, a US Postal Money Order, which is not legally currency yet bears many of the same properties: a printed “face value”, anti-forgery document features, and clear monetary value. Yet, a Postal Money Order is not currency, it is only a document used by the Postal Service to facilitate money transmittal by providing evidence of the deposit of currency by the sender.

The use of a physical document to facilitate money transmission is not required, and “wire” money transmitters such as Western Union make use of an electronic message-passing network to replace a paper money order. The similarity to Bitcoin is more clear in this case: under FinCEN’s interpretation, a person can purchase Bitcoin in USD, and then that Bitcoin can be exchanged for USD later at a different date by a different person, just as a Western Union money order can be purchased and then exchanged for USD by a different person at a different location.

The application of AML regulations to the various parties involved in the Bitcoin economy can be complex. The 2013 guidance states that a person sending or receiving CVCs in exchange for real or virtual goods is not a money transmitter and not subject to FinCEN regulation. However:

A person that creates units of this convertible virtual currency and uses it to purchase real or virtual goods and services is a user of the convertible virtual currency and not subject to regulation as a money transmitter. By contrast, a person that creates units of convertible virtual currency and sells those units to another person for real currency or its equivalent is engaged in transmission to another location and is a money transmitter. In addition, a person is an exchanger and a money transmitter if the person accepts such de-centralized convertible virtual currency from one person and transmits it to another person as part of the acceptance and transfer of currency, funds, or other value that substitutes for currency.

3.5.2 Bitcoin Business Models

Many of the specific use-cases of Bitcoin remained, to some extent, in a gray area in light of this guidance. In ?) further guidance was issued on the applicability of AML regulations to services which allowed for the trading of Bitcoin. This document further emphasizes the nature of CVCs as subject to regulation:

The label applied to any particular type of CVC (such as “digital currency,” “cryptocurrency,” “cryptoasset,” “digital asset,” etc.) is not dispositive of its regulatory treatment under the BSA. Similarly, as money transmission involves the acceptance and transmission of value that substitutes for currency by any means, transactions de-

nominated in CVC will be subject to FinCEN regulations regardless of whether the CVC is represented by a physical or digital token, whether the type of ledger used to record the transactions is centralized or distributed, or the type of technology utilized for the transmission of value.

The 2019 guidance specifically enumerates various cryptocurrency business models and FinCEN’s interpretation of their regulatory position. For example, it states that peer-to-peer exchange services, generally individuals who advertise that they are willing to buy or sell bitcoin, are subject to regulation as money transmitters. It states that “Bitcoin ATMs” are subject to regulation because, unlike conventional ATMs, they do not merely provide access to an account at an institution which is already regulated.

Further, the 2019 guidance clarifies regulation of Bitcoin exchanges. While it does lay out an exception for the case where the actual transaction occurs entirely independently of the platform, it makes clear that in the vast majority of cases (any case where the exchange is involved in facilitating the actual transaction and not merely exchanging bids) exchanges are considered money transmitters and are subject to regulation.

Perhaps most interestingly, anonymity services, including both tumbling/mixing systems and cryptocurrencies specifically designed for anonymity, are specifically called out in the guidance. FinCEN refers to mixers or tumblers as “anonymizing service providers” and there is some nuance in the issue which is explained in the guidance. FinCEN’s position is best summarized this way:

- Vendors of tools for anonymizing currency are not money transmitters subject to regulation if they are completely uninvolved in the actual process—i.e. if they simply distribute a software package which can be used for tumbling coins. In this case they are viewed as a software vendor engaged in that trade and not as a money services business.
- Entities using a tumbling or mixing tool for *their own purposes*, such as protecting their privacy while paying for a physical or virtual good, are “users” of money, not money transmitters, and are not subject to regulation.
- Entities anonymizing currency for *other than their own purposes*, i.e. for customers, are money transmitters and subject to regulation.

More specifically, FinCEN considers the transmission of currency to be for one’s own purposes, and therefore a “use of money” rather than “money transmission”, only if (a) the transmission is carried out only in the process of conducting another type of business, and (b) the transmission of currency is integral to that business, or in other words, it is not reasonable or practical to conduct that business without the specific transmission of currency in question. As a result of the second rule, this exception to money transmitter regulation is often referred to as the *integral exception*.

FinCEN applies the same reasoning to those who develop not just tools but cryptocurrencies (altcoins) with the purpose of protecting anonymity: developing a software program does not constitute money transmission. However, designing and then *selling* an altcoin designed to preserve privacy would make an entity a money transmitter subject to regulation due to their acting as a service provider to move money for others.

FinCEN addresses distributed organizations or distributed applications with the same logic. First, they consider a distributed application, despite having no individual administrator, to be fundamentally of the same nature as a Bitcoin ATM—that is, an automated, mechanical device, which is nonetheless subject to regulation. This would require that AML measures be “designed in” to a distributed app. And, once again, development of the distributed app is not itself a regulated

activity, but engaging in business with the distributed app (e.g. by accepting payments into it) would be subject to regulation.

Finally, the 2019 guidance briefly refers to an important element of FinCEN regulation: “funds travel rule”. This element of BSA regulation requires that financial institutions forwarding currency on to other institutions also forward sufficient information for the next institution to operate their own AML/KYC program [?]. The funds travel rule does not apply to all types of transactions, since it specifically refers to the situation where a financial institution transmits funds to another financial institution (the definition of financial institution used here includes money transmitters, the category under which the strong majority of cryptocurrency businesses fall).

Transactions between a services provider and a customer are not subject to this rule since the entity on one end of the transaction is not a financial institution under the rule. However, the funds travel rule does mean that in the specific case of funds transmitted from one money transmitter to another, it is specifically in violation of regulation to anonymize or pseudonymize the origin of the funds.

3.5.3 Enforcement

FinCEN has become active in the enforcement of AML regulations in the cryptocurrency industry. In April of 2019, FinCEN fined an individual \$35,000 for failure to register as a money transmitter and failing to submit required CTRs. The individual was acting as a Bitcoin dealer, advertising his ability to sell bitcoin via peer-to-peer exchange platforms [?].

FinCEN is additionally pursuing criminal charges against the operator of the exchange BTC-e for similar failures to comply with regulation. BTC-e was operated from Cyprus and Seychelles but considered subject to US regulation due to their business with US customers [?].

FinCEN’s recent aggressive approach to regulation of cryptocurrency may be explained by the extent to which it now factors into potential crime. [?] reports that FinCEN receives over 1,500 SARs from financial institutions each month related to cryptocurrency activity. FinCEN’s lack of access to information on transactions inside of cryptocurrency networks presents a significant possibility that these reports relate to unidentified money laundering.

APPENDIX 4. BITCOIN LAUNDERING AND ANTI-LAUNDERING

4.1 Customer Identification and Mixing

While Bitcoin is widely perceived as being an anonymous system, because of the nature of the distributed blockchain it is in actuality *pseudonymous*. All transactions are visible to the public and can readily be analyzed, but the participants in those transactions are identified by cryptographic keys which are not necessarily tied to any individual identity.

The major focus of Bitcoin KYC efforts has been on identification of customers at the edge. This model, consistent on the surface with FinCEN’s approach to Bitcoin as a convertible virtual currency (that is, a substitute for conventional currency), involves the identification of customers whenever they exchange conventional currency for cryptocurrency, or vice versa.

This type of enforcement has become fairly common. By 2013, now infamous bitcoin exchange Mt. Gox was requiring their customers to provide government-issued identification before buying or selling Bitcoin [?]. This is clearly considered a “bare minimum” component of anti-money laundering, as evidenced by the major role of BTC-e’s failure to require customer identification in the current legal proceedings against them [?].

This approach to KYC is consistent with that used for the money transmission industry, on which FinCEN has generally modeled AML regulation of cryptocurrencies. Money transmitters are required to collect reliable identification for their customers in order to facilitate submission of CTRs and maintenance of an MIL. To control fraud more generally, many money transmitters collect identification on all customers. This is important since FinCEN regulation may require reporting on total amounts of as low as \$1,000 within a single day—a requirement far easier to meet if identification is collected up-front instead of as part of the transaction which “tips the scales” to require recording or reporting.

However, this model falters when applied to Bitcoin. Bitcoin has several properties which are not typical of conventional money transmitters. Most obviously, while not all money transmitters require that the recipient be identified at the time of purchase, secondary trade in money orders is fairly unusual. That is, it would be very unusual to receive a money order in payment and then use that same money order to pay another individual. Instead, money orders are typically exchanged for currency at the earliest convenient opportunity—often at the location where they are received.

There are at least two reasons that this is the case. First, secondary trade in money orders will frequently involve suspicious or fraudulent activity, including money laundering, and complicates the AML requirements that money transmitters face. It greatly simplifies recording of suspicious and large transactions if the sender and recipient are both known at the time the money order is purchased.

Second, money orders cannot readily be fractionated. That is, it is not practical to exchange a money order for \$100 for two money orders of \$50, as is the case with currency. This makes it extremely inconvenient to use money orders for third-party payments, since they are unlikely to be in the correct amount.

Bitcoin presents a very different situation. In fact, the vast majority of users consider Bitcoin to be not a “convertible currency” but a currency itself, and so trade in Bitcoin is the main purpose

of the system, not an unusual edge-case usage. Bitcoin can readily be fractionated, and in fact the complex multiple-input and multiple-output transactions which Bitcoin allows enable more complex division of Bitcoin than is even theoretically possible with money orders.

The result is a significant challenge to customer identification-based AML programs and KYC more generally: Bitcoin cannot effectively be monitored for money laundering by identifying transactions at the edge, because relationship between those transactions may be intractable.

While in theory there is always a connection between a purchase of Bitcoin (using a conventional currency) and later sale of Bitcoin, the nature of the Bitcoin ledger means that this relationship can, intentionally or by natural result of transactions over time, be completely obfuscated. This breaks the chain of “money in” to “money out” and allows Bitcoin to be used as a medium for money laundering, much like the ill-fated Liberty Reserve.

These same mechanisms apply to laundering not only of USD but of Bitcoin itself, since there are various mechanisms by which Bitcoin can be obtained and spent directly. However, from a regulatory perspective, Bitcoin is currently always viewed as a money laundering concern because of its convertibility. This does not mean that AML requirements do not apply to Bitcoin transactions which do not involve conventional currencies, since Bitcoin is convertible whether or not it is actually converted.

This obfuscation of the relationship between Bitcoin inputs and outputs can occur naturally over time if the Bitcoin passes through multiple hands, but money laundering generally occurs when a single individual or an organization has the need to maintain possession of Bitcoin they already have without it being apparent where it came from. Since transactions which are not spent will not be divided and recombined, in this case the Bitcoin will remain very easy to trace if artificial measures are not taken to obfuscate its source. This artificial obfuscation is exactly the service provided by bitcoin tumblers and mixers.

4.2 Tumbling and Mixing In Practice

Bitcoin mixers, tumblers, or laundries are all largely equivalent terms for services which somehow break the link between Bitcoin submitted by a user and Bitcoin eventually received by that same user (but to a different address). This effectively makes the Bitcoin untraceable, and for this service mixers charge a fee.

Bitcoin mixers fall into several categories depending on their architecture. Early mixers were simple centralized systems, and this approach is still common. Several mixers are now available based on complex Bitcoin transaction types such as CoinJoin. Finally, some mixers now exist which always pay out using newly mined coins, or use a “chip” payment approach.

4.2.1 Centralized Mixers

A centralized Bitcoin mixer is the simplest type to analyze. In a centralized mixer, a user submits some amount of Bitcoin to an address provided by the mixing service. The mixing service then creates a large number of transactions over a period of time, using multiple inputs and outputs to mix all of the Bitcoin received together, preventing easy analysis of the flow of Bitcoin.

At some point later, the centralized mixing service adds the user’s receiving address or addresses as outputs to one or more transactions, paying the user out the same amount they put in, minus a mixing fee.

Centralized mixers are relatively easy to design because they make use of a central coordinating authority which tracks the correlation between all inputs and outputs, allowing easy tracking to en-

sure that each user receives their money. This comes with the significant downside that centralized mixers are at least temporarily aware of the mapping of inputs to outputs, and this information could be stolen or otherwise obtained, for example by law enforcement seizure of the service. To mitigate this possibility, centralized mixers generally promise to retain no records of mixing after the final output transaction to each customer, although in practice customers have limited assurance of this promise.

Considering the operations of a central mixer makes several of the challenges in Bitcoin mixing clear. If we wanted to analyze the operation of a central mixer in order to re-correlate the inputs and outputs, an obvious approach would be based on amounts: any inbound transaction should have one or more outbound transactions which total to the same amount minus the mixing fee. Centralized mixers commonly use several approaches to prevent matching up inputs and outputs in this fashion:

- The mixing fee is typically randomized or user-selectable within a range, which prevents the exact total of the outputs being calculated based on an input.
- The output transactions are usually split across multiple independent addresses controlled by the user, so that no one address receives payments with the correct total. However, in some cases users may re-consolidate their funds afterwards, allowing analysis on this basis.
- The output transactions are typically delayed by a random period of time, often distributed around 24 hours. This long and randomized delay makes correlation based on timing more difficult.

While the centralized mixing service receives funds and coordinates paying funds out to users, the coins under control of the service at any given time are continuously “mixed” by making frequent transactions of multiple inputs and outputs between different wallets controlled by the service. The purpose of this mixing is to make it difficult to identify all of the wallets controlled by the mixing service and the flow of funds between them, allowing each customer transaction to “disappear into the noise” without the need for the mixing service to maintain very large reserves.

These techniques, combined with a large customer volume, are usually enough to make analysis difficult but not impossible. [1] investigates the operation of several popular mixing services at the time (most of which are still operational today) and finds that the structure of their mixing process and output transactions is fairly simple.

Despite the mixing process being intended to prevent simple discovery of all input and output transactions, many services are found to rely on transactions in and out of a central wallet, construct transactions with funds just received from customers as each input, or otherwise behave in ways that allow for discovery of the addresses belonging to the service and its customers.

A common pattern used for output addresses is that of a “peeling chain,” in which a wallet is funded with some amount of Bitcoin and then sends some of its value to a customer and some to a newly created wallet. That second wallet then sends some to a customer and some to a newly created wallet. This pattern continues until the original funds are exhausted. While the peeling chain structure prevents having multiple customers as outputs on a single transaction, it is still vulnerable to trivial graph analysis.

[1] attempts a further analysis of centralized mixing services with, to quote the title, “mixed results.” Centralized mixing services appear to have improved in their design by the time of that paper (for the two that worked at all), but the blockchain analysis is fairly limited.

It seems surprisingly how limited experimental investigation of the operation of centralized mixers has been. This can likely be attributed to two causes: first, it can become expensive to

obtain enough Bitcoin to evaluate multiple mixing services. Second, the academic community already considers centralized mixers to be of poor quality due to their vulnerability to compromise or seizure of the central service. Instead, research focuses primarily on decentralized mixing services.

4.2.2 CoinJoin and other multi-party transactions

More recently, mixers have started to emerge which are based on multi-signature or multisig transactions. These transactions require multiple signatures verifying their inputs, and by use of coordinating software can use inputs from multiple independent users. In this case a signature is required from every user to make the transaction valid, creating a kind of escrow system.

Such a system was prominently proposed by Bitcoin Forums user [?\)](#) under the name CoinJoin. CoinJoin describes a method where a set of users mutually interested in improving their anonymity each agree to sign a transaction which takes inputs from each of them and provides outputs to new addresses under the control of each member of the transaction.

CoinJoin eliminates the need for a mixer with an ongoing internal mixing process: instead of transactions being made unlinkable by a circuitous route between them, CoinJoin uses only a single transaction, but prevents any documentation on the blockchain of the relationship between the inputs and outputs.

CoinJoin is not without limitations. First, because only a single transaction is used, it is fairly easy to enumerate all of the output addresses which could possibly be associated with mixing via CoinJoin. This invites any of a variety of techniques to re-associate output addresses with a user attempting to anonymize their currency. To prevent trivial correlation, the amounts mixed must either be the same for all users, or the inputs must be distributed to multiple outputs each (potentially a large number of outputs if the total transaction is small).

Second, CoinJoin does not address the question of coordination to create multisignature transactions. A group of users must find each other and agree on the contents of the transaction, including the outputs. Even if this process is anonymized using a system such as Tor, each user will need to propose the output address they wish to receive their mixed coins to, and so the mapping of inputs to outputs is at least temporarily known to participants in the mixing.

Despite these shortcomings, CoinJoin has the significant advantage of being largely immune to fraud. Every party involved in the transaction must sign it for it to be valid, and these signatures cover the outputs as well. There are few options that provide a user with the same level of assurance that they control where their money is sent.

Further, while not widely implemented, the Mixcoin system in [?\)](#) uses a verifiable mix network architecture and addresses several of the shortcomings of CoinJoin, offering a larger anonymity set and a transaction setup process which is less vulnerable to malicious disruption.

As will be discussed later, these systems have received little actual usage. However, they form one part of a large body of research into more fully anonymous methods of mixing.

4.2.3 Non-Tainting Mixers

A major concern in the privacy of a Bitcoin mixing service is its vulnerability to *taint analysis*. Taint analysis is an obvious form of reidentification which relies on tracing coins through a set of transactions, and will be discussed in more depth later. Centralized Bitcoin mixers typically prevent taint analysis by ensuring that input coins are never associated with the outputs to the same user. However, for a complex mixing system or a mixer with few users, this can be difficult to achieve and is a common source of defects.

Two methods have been developed which are immune to taint analysis, both based on the concept of paying users out with coins that could not possibly have been associated with the inputs. First, a mixer could maintain a reserve of newly mined coins (that is, coins created by the first transaction of each block based on mining fees) and use those coins to pay out users. Because they are newly created, they could not possibly be associated with the inputs.

Second, a mixer can pre-fund outputs before a mixing transaction is started. This essentially reverses time for the transaction, so that the outputs are paid before the inputs are received and so the two cannot possibly be associated with each other. In practice, this means creating multiple wallets, depositing varying balances to them, and then paying out users by revealing to them the private keys for a set of wallets adding up to the value being mixed.

This has the added advantage of giving the user complete control over when the mixed coins appear to leave the mixing service. However, because the mixing service must have possessed the private keys for those wallets when they were generated, they may be subject to theft until the user transfers their value to a wallet under only the user's control. This incentive to promptly remove the value from the payout wallets significantly reduces the advantage of placing payout under user control.

While both of these methods prevent taint analysis, they leave the mixing system open to reidentification of users based on correlation of funds into and out of the mixing service, just like a traditional centralized mixer. Neither of these methods are widespread.

4.3 Wallet Re-Identification in Practice

Overall privacy in the Bitcoin system is notably treated by [?]. While Bitcoin is pseudonymous, the open nature of the distributed ledger invites a variety of methods to associate wallets with users and identify those users. This work makes it clear that the pseudonymity of Bitcoin is not sufficient to protect user privacy and that reasonably modest effort and access to data allows for the recovery of significant information about user activity from the blockchain.

[?] describes a tool that performs automated annotation on the Bitcoin blockchain to match wallets to known users obtained from public data sets (e.g. users disclosing their own addresses on the internet in order to receive payments or donations). This tool is successful in both reidentifying a portion of users and automatically detecting unusual or suspicious activity including significant real-world events.

[?] conducts an experiment using a “transaction simulator” emulating a retail market environment with a well-known set of users (specifically that of a college campus) and develops a method which successfully reidentifies 40% of addresses.

Methods of tracing the flow of currency specifically through Bitcoin mixers have not been widely published. However, from informal discussions and occasional academic analysis [e.g. [?], [?]], several methods are known.

First is taint analysis. In taint analysis, coins belonging to a suspect individual are “tainted” or annotated within the blockchain. Those coins are then traced through transactions, with all of the outputs of any transaction receiving tainted coins as inputs becoming tainted as well. This creates an expanding set of tainted coins, each of which is potentially still a payment to or by the suspected user.

The set of tainted coins can often be reduced by relying on user traits. For example, a user mixing Bitcoin may use the mixed coins to make a purchase requiring a sum greater than each individual transaction. As a result, they are likely to combine more than one of the outputs from mixing into either a single wallet or as inputs to a single transaction.

When performing taint analysis, the analyst can focus on only those outputs which are combined at some point in the future. This can significantly reduce the scope of the tainted coins and still presents a good chance of locating the mixed bitcoin if time is allowed for the user to begin to use their mixed coins.

Taint analysis can also be used in reverse, by starting from output transactions and tainting the inputs of those transactions, in that way working back to possible inputs from the original user.

Taint analysis is less effective today because many mixers have been architected to avoid coins tainted by a mixing input ever reaching any outputs of the same transaction. There are several methods to achieve this isolation, but the simplest is to simply maintain multiple isolated mixing pools and simply pay out from a different pool than the user paid into.

Second, the inputs and outputs of Bitcoin mixing can often be located by simple correlation of transaction sums. While mixers usually delay their outgoing payments by a randomized period of time and break up the output into multiple transactions of random amounts, users will expect their mixed Bitcoin to be available to them within a reasonable period of time, and so for any input to the mixer outputs of nearly the same sum can be expected to go to wallets controlled by the user within the next several days.

Analyzing mixers by locating the outputs corresponding to inputs requires identifying the entire set of addresses controlled by the mixer, so that all outputs of the mixer in the relevant time period can be identified. Mixers try to avoid this behavior, but [1] and [2] show that simple clustering is effective in identifying the pools used by most mixers at the time.

CoinJoin and other multi-signature transaction based mixing systems make taint analysis largely irrelevant, as they intentionally operate in such a way that every output of each mix is tainted by every input, resulting in a large anonymity set even after taint analysis. However, CoinJoin-based mixers are still subject to correlation of amounts and analysis of later use of mixed coins for behavioral cues.

Finally, it is an important property of mixing services that the anonymity set of mixer users is generally limited by the number of users of the mixer. This creates an incentive to use a highly popular mixing service while making mixers that are new or more technically complex to use less valuable. This may help to explain why adoption of more sophisticated types of mixing has been slow.

This discussion has focused only on methods of reidentifying mixed Bitcoins based on analysis of transactions in the blockchain. In practice, it may often be more practical to identify the history or later use of mixed coins by non-blockchain methods, such as by exploiting user behavior and other attacks on internet anonymity (e.g. tracking of internet traffic).

4.4 User Behavior and Bitcoin Implementation

As with many cryptosystems, the greatest weakness in Bitcoin anonymity may be not the system itself but user behavior and implementation details.

User behavior presents many challenges. [3] show that many users of the Bitcoin system can be identified by their presence in public data sets of Bitcoin addresses, and suggests that services such as exchanges which interact with a large number of users are positioned to potentially reidentify a large portion of active users through their own identity records and then association of users based on transactions.

Similarly, the same work found that in cases of specific interest, such as a high-profile theft of a large sum of Bitcoin, significant information can be learned about the disposition of that money and the thief's relation to other events (including other thefts) based on simple review of

the flow of money. This information would provide a significant starting point to a law enforcement investigation.

Bitcoin mixers are largely intended to mitigate these well-known limitations in Bitcoin privacy, but they are effective only in preventing or complicating blockchain-based analysis. A variety of non-blockchain methods can be used to identify users.

?) discussed the implementation challenges to privacy related to the P2P protocol used to distribute transactions through the Bitcoin system prior to creation of each block. Because a node which wishes to introduce a transaction to the blockchain must widely broadcast that transaction, in most cases the first IP address which any given node receives a transaction from represents the computer which first introduced that transaction. In the case of typical retail purchases, this could present a very simple way of identifying the purchaser despite mixing.

These details of the P2P distribution network are not recorded permanently. However, it is a fairly simple matter to modify a Bitcoin client to record this information (or make use of a separate packet sniffer and protocol analyzer in front of an unmodified Bitcoin client) and to record the "first seen" origins of transactions over a long term. Doing so from multiple devices at different positions in the internet would increase the probability of successfully receiving a transaction directly from the user that created it.

Similar problems are presented by the non-Bitcoin aspects of most Bitcoin transactions. Purchases of goods with Bitcoin, exchange of Bitcoin for other currencies, and even mixing of Bitcoin are almost always arranged by use of a website. A user may reveal their IP address in visiting such a website, and especially in the case of purchases, may need to reveal other identifying information such as a shipping address.

?) suggests that users are aware of these concerns and take precautionary measures such as the use of the Tor anonymizing network and fictitious personal information. However, purchase of physical goods and exchange to conventional currencies both pose the challenge of the product needing to be delivered to the buyer, which almost always introduces an opportunity to learn more about the buyer.

Further, it is well-known that users of the Tor system often compromise their privacy through user behavior, limiting the protection that this system offers [e.g. ?), ?)]. Ironically, ?) found that Tor users can sometimes be reidentified based on their Bitcoin activity. Certainly the combination of the two do not provide perfect protection.

Bitcoin mixers are particularly significant in this landscape because they may be ideally positioned to collect information about users involved in illegal or suspect activity. Discussions on the Bitcoin forums show that users are well aware that mixing services could be "honey pots" operated by law enforcement or intelligence agencies to gather information on Bitcoin users. The anonymous nature of mixing services makes it difficult to rule this possibility out for any given service.

APPENDIX 5. MIXING SERVICES AND AML SERVICES

5.1 Review of Bitcoin Mixing Services

To better understand the ecosystem of Bitcoin mixing services, an attempt was made to enumerate mixing services in use. The Bitcoin Forums are an extremely popular central resource on the Bitcoin Community, and a large portion of all Bitcoin services, particularly those oriented towards every-day users and those concerned about privacy, are announced and promoted there.

The “Services and Announcements” section was reviewed based on keyword search and manual discovery to find all announcements of Bitcoin mixing services. Further, discussions on privacy and mixing were reviewed for any mixing services that they mentioned. The result is an enumeration of 69 Bitcoin mixing services that were either in use or announced at some point in time. The full list is included as Appendix A.

Examination of a large number of Bitcoin mixing services shows the difficulty of establishing commercial trust in a rapidly-changing marketplace with significant use of privacy technology. Many mixers have very similar names, which may be an intentional misrepresentation or simply a result of the desire for obvious names (e.g. “bitmixer”). Many mixers operate Tor hidden services, which have randomized addresses and so are often duplicated by fraudulent operators—and may also legitimately change their addresses.

5.1.1 Marketing

It is difficult to overstate the significance of the Bitcoin Forums to the Bitcoin community and economy. These forums serve as an almost universally known central point for discussion and advertising of Bitcoin services. As a result, mixing services are often announced first on the Bitcoin Forums and later advertised there.

Following the announcement of a service (where the features are usually listed, as well as clearnet and Tor addresses), mixing services may take a number of further steps to promote themselves and gain trust. First, some new services will agree to place a small number of coins in “escrow” with a well-established member of the Bitcoin Forums community. This ad-hoc arrangement gives users some confidence that they will not lose their Bitcoin, as if the mixing service fails to return their coins they can use the signed “proof of mixing” letter that these services generally issue to request that the trusted user make them whole.

Second, services may use paid advertising. Because of the legal issues surrounding these services it is not common for them to advertise by conventional means. Instead, they may launch a “signature campaign” wherein they pay prominent Bitcoin Forums members to paste advertisements for the service into their forum post signatures. Services may also promote themselves by purchasing advertising on Tor hidden services. In both cases, payment can be made in the form of Bitcoin.

5.1.2 Features

Newly launched laundering services are entering a crowded market. Users are very hesitant to trust new services, and with several well-established services operating at almost every point in Bitcoin’s history it is difficult for new services to find users.

New mixers use several methods to attract new users. The first, and perhaps most important, is the sets of claims made by new services when they are announced. By far the most common marketing claim made by new mixers is “zero taint.” This indicates that users, or at least mixer operators, are well aware of taint analysis as an analysis technique. However, mixing services virtually never advertise that they prevent *other* mixing methods, suggesting that the community is focused primarily on evading taint analysis at the cost of preventing other methods of analysis.

Users of mixing services appear to be aware of the fact that centralized mixing services are capable of reidentifying their users. As a result, many mixing services explicitly advertise that they do not retain any logs or other identifying information on users. It is also common for services to advertise that they do not require users to “sign up” for an account or provide any identifying information (such as an email address), although many mixing services do require that users set up an account and provide additional identifying information.

Another common advertising claim is that user of the mixer does not require Javascript. This is in response to the large number of users who access mixing services via the Tor Browser Bundle, which in many versions blocks Javascript execution by default.

Common features advertised by mixing services include:

- Random delay times from input to output, which prevent simple correlation analysis based on timing.
- Randomized fees, which make correlation analysis based on amount more difficult.
- Cryptographically signed documents stating the mixing service’s obligation to pay out, which could be used to prove that the mixer fraudulently failed to return a user’s coins.

These features show user awareness of the risks of using centralized mixing services, including both fraud and capture of their information from the service operator.

It is important to note that the two common features intended to frustrate correlation analysis, random delay times and randomized fees, are of only limited value if an analyst is able to determine the full set of wallets in use by a mixing service. However, mixing services almost never make any claims about the quality or complexity of their internal mixing process. This likely reflects primarily the difficulty of describing these internal mixing processes succinctly, even though they appear to generally be fairly simple.

5.1.3 Failure to Launch

69 mixing services over the relatively short lifetime of Bitcoin seems like a high rate of new service introductions. Were they evenly distributed, this would be about eight new services a year, but in practice the rate of launch of new mixing services has increased, with more than once introduced per month starting in 2017.

This large set of new mixing services are seldom successful. Of the mixing services identified, 26 or 38% simply failed to generate any interest. Few of these are still available, suggesting that they did little business and so their operators closed them down.

Many announcements generate one or two replies suggesting that there is no reason to trust a new service based on a brief announcement. This suggests that the problem of fraudulent services

(which simply keep deposits) presents a real challenge to adoption of new mixing services, and so the community will prefer to continue to use a small number of trusted services.

The tendency of users to distrust new services poses an additional challenge to the adoption of improved mixing technology, since even for services which employ technical means to prevent fraud such as CoinJoin, user may be hesitant to leave services which they know to be reliable.

5.1.4 Scams

One of the most obvious conclusions from an enumeration of announced Bitcoin mixing services is how few have survived to the present day. In fact, the majority of mixing services announced appear to have been scams.

The problem of Bitcoin services, particularly those which advertise anonymity, being fraudulent is well known. The Bitcoin Forums feature a thread reputation system in an attempt to mitigate fraud, with users either “vouching” for or “warning” of services based on their knowledge and experiences.

In experimenting with Bitcoin mixers, ?) lost their Bitcoin to a scam in three out of five attempts. ?) eliminated a possible scam from their planned set of mixing services. And with Bitcoin services with profiles as large as Mt. Gox disappearing and taking their users funds with them, fraud is obviously a possibility with mixing services.

In this enumeration, mixers were considered a likely scam if there were indications in the Bitcoin Forums discussion around them that they were a scam, or if they were reported later elsewhere on the internet to have disappeared with user funds. It is likely that this undercounts actual scams since many mixers failed to generate any attention at all. It also undercounts scams occurring via Tor hidden services, where the same service will often frequently generate new addresses, making counting difficult.

With this likely undercount in mind, 19, or 28%, of discovered Bitcoin mixers were likely scams. This set of fraudulent mixing services reveals an interesting challenge in the loosely organized and often informal Bitcoin industry: many simply duplicated the name of an existing, reputable mixer in order to gather coins from confused customers arriving at the wrong website.

A notable example surrounds Helix Light by Grams, a well-known mixing service from the popular darknet search engine Grams. Helix Light was discontinued by its operator, but a visually identical service also calling itself Helix Light appeared at a different address and solicited payments. This seems to have been a simple scam on users who had not heard that Helix Light had shut down.

The problem of evil twins is further compounded by the heavy use of Tor hidden services, which have long, randomized addresses which are often not very memorable or recognizable. This makes it easy to post an address on the internet which claims to be for a well-known service but actually leads to a fraudulent duplicate.

The dual internet/Tor nature of many mixers opens up an interesting novel class of evil twin. Bitcoin Fog is a popular and trusted mixing service which operates only as a Tor hidden service. The visually identical Bitcoin Fog website on the public internet was fraudulently created by a different operator, and capitalizes on Bitcoin Fog’s well-known name and lack of internet presence.

5.1.5 Types of Mixers

Of the Bitcoin mixers considered, the vast majority are centralized mixers which make no particular claims about their architecture. 62, or 90% of announced mixers are centralized mixers. Eliminating apparently fraudulent services and services which failed to gain interest, 77% are still centralized.

This result is surprising considering the disadvantages of centralized mixers, which include a greater potential for fraud, vulnerability to seizure or exfiltration of records by law enforcement, and opaque operations which may obscure poor design. However, it is understandable in consideration of usability factors.

Centralized mixers are generally highly user-friendly. A user need only access a website (possibly through the Tor network), enter some information, provide output addresses and deposit coins to an input address from the wallet of their choice, potentially including a custodial wallet service (which operates a Bitcoin client and manages a user’s keys on behalf of that user for ease of use).

This type of service is easy to understand and use for novice users, and are highly compatible with whatever wallet solution a user is already familiar with. They are also easy to discover via search engines, darknet search engines, and websites where Bitcoin is discussed.

On the other hand, decentralized mixing systems require that the user download and use a client software program. Depending on the system, the user may also need to use the client software program as their wallet, importing or generating new keypairs. This complexity, and the risk of using an unknown application, is likely a discouragement to new users.

5.1.6 Anonymity Protections

To prevent mixing services (or others) identifying users based on their IP addresses, users are commonly advised to access mixers using an anonymizing network such as Tor. To facilitate this, many mixers provide a Tor hidden service, and some operate exclusively in this fashion.

The function of Tor hidden services is to provide anonymity of the *provider* of the service, and so offering a hidden service alongside a “clearnet” or public internet website negates this security [?]. However, this is a common practice. The willingness of mixer operators to shed anonymity by providing a clearnet website (with associated domain registration, IP address, and other information which could be used to identify them) suggests that many users are not sophisticated enough or not willing to access such services through Tor.

On the other hand, the decision of these mixers to provide a hidden service despite also offering a clearnet service, when this has a reduced advantage to their users, suggests pressure to appear to use various conventionally expected security measures.

Of 69 hidden services, 39, more than, provide both a clearnet website and a hidden service. 8 are accessible exclusively via a Tor hidden service. 19 are accessible exclusively via clearnet. The remainder are client-based mixers which do not use their website as part of mixer operation.

5.1.7 Operator and Law Enforcement Actions

Of the 69 services examined, only one was publicly shut down by law enforcement. This service, bestmixer.io, was seized by the Dutch Financial Criminal Investigative Service after a multinational investigation. The primary fault in bestmixer’s efforts to evade law enforcement seems to have simply been their location in the European Union—a vulnerability they were apparently aware of as, according to (?), they falsely advertised their location as Curacao.

There are no other reported incidents of law enforcement seizure of bitcoin mixing services, suggesting that global enforcement of AML regulations and criminal laws is limited. While it is difficult to determine the country of origin of mixing services, a likely explanation for limited law enforcement is that the most successful mixing services are located in countries with little or no AML policy.

Far more common than interference by law enforcement is the decision of an operator to shut down their own mixing service. At least 11 services were closed by their operator, some after

they had failed to gain any significant attention, but others were shut down despite widespread popularity.

The major Bitcoin mixer bitmixer.io closed doors after the operator posted online that they had realized that “Bitcoin is transparent non-anonymous system by design”¹. Although it is difficult to verify this claim, bitmixer.io has stated that they processed 65,000 BTC per month, making it a very significant player in the market and a surprising service to so abruptly close [?].

5.2 Observations

5.2.1 Difficult to Establish

While the number of Bitcoin mixing services which have existed is fairly large (with 69 almost certainly being an undercount), the number of services which are trusted by the broader Bitcoin community is fairly small—perhaps a half dozen. One often-linked-to list on the Bitcoin Forums includes 14 mixing services not thought to be scams, several of which are still fairly new.

For potential users of Bitcoin mixers, the greatest hazard appears to be common scams, including simple duplicates of popular mixing services. Users seem to also face the challenge of technical complexity, with the most easily used mixers (centralized services available on the clearnet) also being the highest risk for both fraud and reidentification.

The greatest hazard faced by mixing services themselves seems to be a crowded marketplace in which trust is difficult to earn. Mixing services far more often fail due to the inability to earn customers than due to law enforcement action. While the possibility of criminal charges is no doubt a factor in the decision of some operators to shut down their mixing services, the challenges of remaining profitable while keeping up with changing technology and the instability of the Bitcoin market are likely also significant factors.

5.2.2 Community Response to Scams

Because of the ease of launching a Bitcoin service and the anonymity with which these services often operate, scams are a significant challenge in the Bitcoin grey market.

The community has developed several ad-hoc measures to detect mixing services which are merely scams that intend to shut down and retain the funds which had been sent to them. The first is the practice of some services of placing funds in escrow with respected members of the Bitcoin Forums. This provides the community with assurance that they will not lose their funds, to the extent that they trust the individual who holds the escrowed Bitcoin.

Further, because the Bitcoin Forums are a common central point for information on Bitcoin services, its users have developed an informal reputation system. Users often report any apparently fraudulent behavior on the announcement page for a service, and discuss scams in other sections of the forum. This evolved into a system where users can “vouch” for a service or report a scam. If multiple users accuse an announced service of operating as a scam, a clear warning displays above the thread.

¹<https://bitcointalk.org/index.php?topic=2042470.0>

5.3 Blockchain Analysis Services

In response to the complexities of AML requirements as applied to Bitcoin and the needs of law enforcement and civil litigation, a number of tools and services have emerged which provide blockchain analysis to financial services businesses and investigators.

These can broadly be placed in two categories: AML/KYC products which are intended for use *prior* to any incident as a risk management process, and investigatory products intended for use *after* an incident. These two classes of tools approach the problem from somewhat different directions. KYC products generally attempt to determine whether or not a transaction has been intentionally anonymized, and if so do little further than indicate a high risk. Investigatory products, on the other hand, are often used explicitly because the transaction has been anonymized, and must attempt to reverse the anonymization process.

Many services which advertise themselves as blockchain KYC/AML solutions (or as more general solutions which are applicable to cryptocurrency as well) only address the CIP component by collecting and verifying customer identification documents. While important, this component of AML compliance does not involve the actual analysis of the blockchain and so is excluded from this discussion.

A complete list of analysis services and tools considered is included in appendix B.

5.3.1 AML/KYC Services

Five providers offer a cryptocurrency AML solution based on analysis of the blockchain. It is difficult to provide substantial analysis of these services because they publish very little about their internal methodology, perhaps out of concern that it could be contravened by actors with knowledge of the algorithms in use. This substantial challenge to academic analysis has previously been encountered by e.g. ?). Some inferences about the state of the art in Bitcoin transaction risk analysis can be drawn from the open literature on the subject.

?) discuss two important heuristics in analysis of the blockchain: first, all of the inputs to a transaction generally belong to the same person. Second, there is usually a change output on the transaction which also belongs to the same person as these inputs, subject to certain constraints on the identification of the change address. It is shown that these two heuristics allow substantial clustering of Bitcoin addresses by ownership, and that these clusters allow for reidentification of addresses belonging to users one has interacted with (and out to additional degrees).

In a different vein, ?) introduce a set of methods for graph analysis of money laundering in general (that is, not specifically for cryptocurrency). It is shown that link analysis can be used to identify likely participants in money laundering based on characteristic patterns, such as dividing money across multiple activities and then recombining at a later time.

These approaches can be combined to detect finances involved in money laundering. In fact, this is the easiest formulation of the problem of analyzing Bitcoin mixing services, since the only requirement is to identify the outputs of a mixing service, with no need to identify inputs or their relations.

First, addresses belonging to mixing services are identified. This can be done by a number of methods, but the most obvious is to initiate transactions with mixing services so that they reveal an address to be used as an input.

Clustering methods are then used to identify further addresses related to a mixing service. ?) and ?) find that addresses in use by mixing services tend to be tightly clustered according to

well-known heuristics such as those in ?), and so a large portion of the addresses used by a mixing service can be identified in this fashion.

The result is a database of addresses known to belong to mixing services. Any transaction can then be traced back in terms of its inputs. Any path back which leads to a mixing service indicates a higher risk of money laundering or fraudulent activity, with that risk decaying according to the number of steps in between and the portion of the transaction funded by apparently laundered coins.

This method is essentially taint analysis performed in reverse: from a given transaction, coins are traced backwards in order to establish whether or not they are tainted.

5.3.2 Investigative Tools

A number of tools also exist which are intended for post-hoc investigation of blockchain activity. These are intended primarily for law enforcement and investigators for litigation, and are oriented around understanding the flow of money that is known to have been involved in a criminal or otherwise suspect act.

These tools usually combine visualization tools along with annotation tools, and may include clustering features. Visualization tools allow an investigator to easily follow the flow of bitcoin between transactions by visually moving between inputs and outputs. Transactions are usually presented in a graph format but the view is often simplified or reduced in scope to maintain ease of use.

Annotation in forensic tools usually consists of a vendor-provided database of addresses annotated by known owners, and the ability for the user to add their own annotations to addresses as their owners are determined. Clustering tools may assist in annotating other addresses which apparently belong to the same owner.

Because investigative tools are intended for manual post-hoc use, they are impractical for use as part of an AML program because of the time and effort which would be required to manually investigate a large number of transactions.

5.3.3 Implications for the Bitcoin Market

Risk scoring tools for Bitcoin transactions are becoming increasingly common as exchanges and other service providers develop AML compliance programs. The result is that an individual possessing Bitcoin which is tainted may have a difficult time using it—a situation similar to counterfeit money, and with the similar disadvantage that the holder of such Bitcoin may have obtained it legitimately and not be aware of its suspect past.

This has interesting implications on the broader Bitcoin market. ?) observed that Bitcoin may have an additional axis of value beyond its denomination: that of trust. Bitcoin which has a suspect or high-risk past may be less valuable to many users than Bitcoin without such a taint. Further, any person receiving Bitcoin may run a risk that it is tainted and is not only less valuable in its own right but may even reduce the value of any Bitcoin with which it is mixed.

?), explores the implications of this heterogeneous aspect of Bitcoin, which could significantly complicate the Bitcoin market by making Bitcoin payments a higher-risk activity which will be less reliable due to the need of parties involved to protect themselves by rejecting high-risk transactions. Considering the likelihood that, over a long span of time, all circulating Bitcoin will pass through a mixing service, this presents a significant practical problem.

5.3.4 Dusting

This problem also presents the possibility of intentional manipulation of Bitcoin risk. Indeed, such an event has occurred at least once. On October 23rd of 2018, major Bitcoin mixing service BestMixer.io sent small amounts of Bitcoin referred to as “dust” to a large number of recipients. While BestMixer.io publicly identified this as a new form of advertising, it is broadly thought to have been an effort to foil automated AML analysis by artificially tainting a large number of legitimate wallets [?].

This activity can be located on the blockchain using relatively simple heuristic analysis, by identifying transactions with large numbers of small outputs. Other such transactions also appear on the blockchain, suggesting that dusting attacks have occurred in multiple cases and perhaps by multiple actors. It is not entirely clear whether these were intended to complicate taint analysis or enable it by generating artificial inputs to known addresses in the hope that they would be spent. However, at the minimum, dusting activity does somewhat complicate taint analysis.

This type of activity has also occurred in the case of Litecoin, although there is some debate around the motivations underlying the dusting [?]. Further, it is not clear if such activity has continued in the Bitcoin blockchain since the shutdown of BestMixer.io by law enforcement.

It is not clear if AML analysis services have taken countermeasures against incorrectly assigning high risk to wallets due to previous dusting attacks. Correcting for this artificial activity is fairly straightforward: since the BestMixer transactions sent amounts between 666 and 888 Satoshi, very small amounts, an AML service would be well advised to simply ignore inputs of such small size unless they cumulatively add to a larger value in a single wallet.

APPENDIX 6. FUTURE DIRECTIONS

While the current landscape of Bitcoin AML is unsettled and carries significant unknowns, the Bitcoin economy continues to undergo rapid development which may reduce the efficacy of AML analysis in the future. At the same time, analysis tools continue to increase in sophistication, driven by the needs of law enforcement and financial regulators.

This section presents an overview of current developments in the academic community and industry which may become important factors in the future.

6.1 Improved Laundering Techniques

6.1.1 Multi-Signature Transactions

The CoinJoin multi-signature transaction approach to mixing introduced by ?) remains an important component of newly introduced Bitcoin mixing methods. CoinJoin offers the possibility of a Bitcoin mix which eliminates the possibility of fraud by the mix operator, which appears to be one of, if not the, greatest practical challenges in anonymizing Bitcoin.

Many improved methods for Bitcoin mixing currently proposed are derived from CoinJoin or otherwise use a single multi-signature mix transaction as the core of their design. However, the dearth of CoinJoin or other multi-signature transaction systems currently in use speaks to the limitations in this design.

First, the original CoinJoin proposal describes only the transaction structure of a CoinJoin mix, but does not propose an anonymous way for the parties of such a transaction to agree on its inputs and outputs. This is analogous to a cryptographic specification that does not specify a key exchange process, and as in that situation, the details of coordinating CoinJoin transactions anonymously are complex and difficult to implement.

This imposes practical limitations on CoinJoin beyond what the transactional structure itself would suggest, such as the need for users in most cases to run specialized software to prepare the transaction instead of simply using a website. Much of the current work in improved Bitcoin mixing is directed at solving exactly this problem, that of *multi-signature mix coordination*.

Second, CoinJoin offers anonymity guarantees which are likely worse than, and at least different from, those of most central mixing services. The *anonymity set* of a mixed Bitcoin transaction (or, really, mixed UTXO ready to be used in a future transaction) is the number of independent inputs to the mix process from which the transaction in question cannot be differentiated.

In the case of a conventional mixing service, the anonymity set of a given use of the mixing service is ideally the number of outputs of the mixing service during the maximum time span allowed between the inputs and outputs (e.g. 48 hours). In practice the anonymity set is generally smaller because central mixing services are less than completely effective in foiling correlation analysis, but the anonymity set is still determined by the number of users of the service over a set time span. For a popular service, this number should be very large.

In the case of a CoinJoin transaction, the anonymity set is constrained to the number of users which participated in that particular transaction. Due to the unpopularity of CoinJoin, the relative complexity of transaction setup, and the desire of users to coordinate a transaction within a

reasonable period of time (without a lengthy wait for additional mix partners), this set is usually small. In the case of the most popular CoinJoin-based mixing system, Wasabi Wallet, (?) reports a transaction with an anonymity set of 100 as the record for largest transaction¹.

These small anonymity sets mean that a series of multiple transactions are likely needed to gain sufficient confidence in the anonymity provided by the mixing process. The need for multiple mixing rounds directly increases the time and complexity required for CoinJoin transactions, further limiting their competitiveness with centralized mixers.

Finally, CoinJoin transactions make especially significant the problem of correlation analysis. Centralized Bitcoin mixers generally increase the complexity of correlation analysis by taking measures to prevent investigators easily determining all of their inputs and outputs. “Secret” outputs which cannot easily be associated with the mixer will be left out of a correlation analysis and effectively prevent this type of re-identification.

CoinJoin transactions are inherently ideal for correlation analysis because all relevant inputs and outputs are immediately visible in the single mix transaction. This makes the small anonymity set a particularly significant problem, and means that in practice it is difficult to achieve significant anonymity with a CoinJoin transaction unless all participants mix a fixed amount [?]. This limitation is a recurring theme in Bitcoin mixing research.

6.1.2 Output Splitting

Correlation analysis can be made more difficult by splitting each input into multiple outputs. Because the outputs can be summed in various combinations to discover sets of outputs that likely correlate with each input, the values of the split outputs must be selected in such a way that there are multiple possible combinations corresponding to each input.

?) describes a simple algorithm which splits larger inputs to a CoinJoin transaction based on the differences between those inputs and other smaller inputs. The result is a set of inputs and outputs which can be combined in many different valid ways, reducing the efficacy of correlation analysis. This method allows an anonymity set of nearly the number of parties to the CoinJoin transaction to be preserved despite the mixing of different amounts, but requires that the transaction be coordinated with total knowledge of the mapping of inputs and outputs—that is, it does not preserve the privacy of mix participants from other mix participants.

The CoinJoin algorithm as described in ?) is currently in use in the mixing service JoinMarket, which, more than a service, is an open-source software tool which can be used to participate in CoinJoin transactions.

6.1.3 Shuffling

To provide participants with protection from other participants, some sort of secure multiparty computation (SMPC) method can be used to allow participants to calculate a set of inputs and outputs without knowledge of the mapping of inputs and outputs contributed by users.

One such prominent system described in ?) is referred to as CoinShuffle. CoinShuffle uses a decryption mix network as first described by ?). In this system, multiple layers of encryption are used to allow output addresses to be shuffled by multiple parties, each of which remains unaware of the complete shuffle order.

¹Further, while SegWit mitigates this concern to some extent, transactions with many signatories are larger and thus more costly in terms of processing fees. This is an additional deterrent to creating large anonymity sets with CoinJoin and is a fundamental limitation of CoinJoin-type transactions. Clever use of cryptography and verification scripts in P2SH transactions may present a route around this limitation.

While CoinShuffle prevents any participant knowing the mapping of outputs provided by other participants, it does not address the matter of correlation analysis and so assumes that every user will mix the same number of coins. This poses a significant practical limitation, since it is difficult to coordinate any meaningful number of people wishing to mix the same value at the same time.

CoinShuffle is not currently in significant use, although at one time an implementation was included in the Mycelium Wallet project called Shufflepuff.

6.1.4 Dining Cryptographers

?) introduced the *dining cryptographer's problem*, an algorithm which allows a group of individuals to set up a communications channel in which any individual can send a signal which the other individuals detect but will not be able to attribute to a sender. This method can be generalized to allow a group of individuals to exchange arbitrary messages without any ability to determine the sender of the message.

The dining cryptographer's network or DC-net has become the basis of many published Bitcoin mixing designs, particularly with later extensions which provide accountability for malicious actors which may disrupt the process. Such an algorithm forms the basis of CoinShuffle++, a further development of CoinShuffle described in ?) which achieves much the same aims as CoinShuffle but using a faster and more reliable process for transaction coordination.

CoinShuffle++ is not currently in use.

6.1.5 Reliance on Anonymity Networks

The use of an anonymizing IP network such as Tor is already a common practice within the Bitcoin privacy community. Some designs for improved mixers have been developed which are specifically reliant on the services of an anonymity network. The most notable example is a brief note included with the original description of CoinJoin in ?): Chaumian CoinJoin.

Chaumian CoinJoin is so called because it makes use of a blind signature scheme described in ?). The signatures are used like so: through an anonymity network, each user submits input information and a blinded output address to a mix coordinator, which the coordinator signs. All users then use the anonymity network to generate a new, unlinkable identity (such as by causing the Tor service to generate new routes) and then submit the signed and unblinded output address to the mixing coordinator. The mixing coordinator is able to verify its own signature and proceed to pay out to the output, but because the full output address was only seen after a change of identity the mix coordinator is not able to correlate the inputs and outputs.

In practice, the mix coordinator actually prepares a multisignature CoinJoin transaction, which allows every involved user to validate the honesty of the mixing coordinator before signing. This relatively simple scheme is surprisingly effective, although it comes with the typical CoinJoin limitations of requiring either a fixed amount or an output splitting scheme and reduced anonymity.

The Chaumian CoinJoin scheme is currently in use as the core of the larger ZeroCoin anonymity protocol (which also includes restrictions on wallet behavior in order to prevent defeating of the mixing by e.g. rejoining split outputs), the underlying protocol for Wasabi Wallet.

6.1.6 Off-Blockchain

For a variety of reasons there has been an increasing emphasis in research on methods of conducting Bitcoin transactions entirely without the use of the blockchain. The limited size and rate of blocks poses an ultimate limit on the rate at which the Bitcoin system can process transactions,

and this has motivated systems for conducting transactions independently of the blockchain and later clearing them in summary on the blockchain for permanent recording—a situation analogous to the credit card network which collects transactions to be cleared between banks in batches.

These off-blockchain systems are also useful for mixing, since they allow for transactions to be created and verified using logic completely different from that used in the Bitcoin system, so long as it can be guaranteed that they will later clear in the form of equivalent Bitcoin transactions.

The TumbleBit system described in [10] requires users to escrow Bitcoin on the blockchain in the form of P2SH transactions, and then uses an independent off-blockchain transaction verification system to pass Bitcoin between users. A coordinating server then posts a summary transaction with the cumulative effects of the off-blockchain activity to the blockchain, which is verified by all participants for correctness. The off-blockchain transaction protocol is designed to preserve complete privacy of users.

Similar work has been proposed in [11] as the Bolt protocol. However, Bolt is not compatible with Bitcoin itself, running instead on the ZeroCash altcoin.

At least two implementations of TumbleBit have seen development work, but neither has gained significant usage. As with other cryptographically strong approaches to mixing, TumbleBit comes with the significant limitation of a fixed mixing denomination.

6.1.7 Altcoins

Much of the academic and commercial research into privacy-preserving cryptocurrencies has turned away from Bitcoin and towards altcoins designed with the explicit purpose of preserving user privacy. Most notably, the ZeroCash altcoin introduced in [12] uses a public ledger which stores only reduced information about transactions. Privacy-sensitive information related to transactions, such as input and output addresses, are held off-blockchain by wallet holders, preventing most types of analysis for reidentification.

Privacy-focused altcoins are intrinsically limited in that they do not attract the attention and user-base of Bitcoin, and at the present time such altcoins are obscure compared to Bitcoin. Nonetheless, as privacy continues to be a significant challenge for Bitcoin users, altcoins designed to avoid these challenges may become popular.

6.1.8 Observations

While the field of Bitcoin anonymization is quite active, academic proposals for cryptographically strong mixing systems have seen little real-world usage. This appears to be attributable to several common limitations in these systems: the requirement for a fixed denomination (or significant compromise on anonymity if variable denominations are used), small anonymity sets relative to popular centralized mixers, and the need to use client software rather than a user-friendly web service.

Moreover, most Bitcoin mixing services in widespread usage appear to be profit-motivated. The path to monetization of a decentralized mixing service is less clear, which may have slowed development of such systems. TumbleBit has been adopted as a component of the Stratis special-purpose altcoin system, but it is unclear if it will gain any popularity in this role.

Nonetheless, off-blockchain transaction systems such as TumbleBit and Bolt likely represent the most promising direction for improved Bitcoin anonymity. Ultimately, though, interest in such fundamental redesign for the purposes of privacy will likely be directed towards privacy-centric altcoins such as ZeroCash rather than ongoing development of Bitcoin anonymizing systems.

6.2 Bitcoin Reidentification

Very little research in the area of countering Bitcoin laundering occurs in the open literature. Most newly developed methods are kept proprietary by service providers which rely on them as a part of their product.

In addition to previously discussed work in [?], [?], and [?], there is some other more recently published literature on the analysis of Bitcoin mixing services.

[?] propose a “de-mixing” algorithm based on the correlation of input and output amounts. Their fairly simple approach is tailored to the mixing service Helix by Grams, which was popular at the time of publication. This algorithm largely serves only to formalize an already well-known technique in the analysis of Bitcoin mixers, which is that for any given input there must be some set of outputs which is closely equal to the input (minus fees).

[?] discusses the challenges of analyzing multisignature-transaction type mixing systems, including some heuristics which can be used to simplify the problem. They are able to re-associate the inputs and outputs on a large portion of transactions using a simple analysis method. While older, this is an interesting reference because, as a whitepaper published by a firm offering Bitcoin services, it provides some insight into industry practices.

The preprint [?] presents an openly published algorithm for risk scoring of Bitcoin wallets called TaintRank. This is likely similar to the methods used by commercial scoring providers, and various methods for propagating risk across the graph are evaluated and shown to provide significantly differing results. This emphasizes the challenge of automated analysis at a large scale.

APPENDIX 7. CONCLUSION

7.1 Bitcoin Laundering

While Bitcoin is widely perceived as being an anonymous system, there are significant limitations to the anonymity it ensures in practice. Because the Bitcoin ledger is subject to public inspection, the system is actually pseudonymous rather than anonymous.

To deter reidentification, Bitcoin laundering services arose and have remained common to this day. While there is ongoing research into effective means of anonymizing Bitcoin, the majority of mixing services in use today are based on a relatively simple centralized design in which one coordinating service receives a deposit and sends outputs to the addresses provided by the recipient, with some sort of internal mixing process used to increase the difficulty of taint analysis and correlation analysis.

The degree of anonymity provided by these mixing services is dubious. Critical evaluation of centralized mixing services has often shown their design to be insufficient to prevent correlation analysis, and well-resources analysts may be able to “de-mix” anonymized coins in a large number of cases.

Future development in Bitcoin mixing services generally focuses on decentralized designs which rely on multisignature transactions and some type of privacy-preserving communication technique to prepare these transactions. While available today, these types of mixers are relatively unpopular and do not appreciably increase the difficulty of blockchain analysis to reidentify outputs.

7.2 Bitcoin Anti-Money Laundering

AML regulations in the United States are broadly understood to apply to Bitcoin money services in the majority of cases. However, compliance with AML regulations on the part of the Bitcoin industry has been slow-coming both because of unwillingness to comply and the intrinsic difficulty of satisfying AML requirements while offering services in a currency explicitly advertised as anonymity-preserving.

AML regulations require money transmitters doing business in Bitcoin to take sufficient measures to identify their customers, prevent facilitating money laundering, and identify and report possible money laundering after the fact.

To satisfy these requirements, many Bitcoin services rely on identification of their customer and monitoring of their transactions involving traditional currencies—a model which can be seen as AML “at the edge,” intended to enforce AML regulations before cryptocurrencies become involved.

However, this model is not applicable in all cases and is viewed as insufficient to prevent the use of Bitcoin for criminal activity. There is a need to perform AML activities for Bitcoin transactions themselves. These activities crucially involve refusal to do business with money launderers.

Several service providers now offer risk scoring services which identify possible money laundering on the blockchain, including the use of mixing services. These scoring services rely on blockchain analysis techniques to identify coins which are the result of likely mixing activity.

The exact methods used by these services are proprietary, but are likely similar in nature to a number of well-known heuristics and analysis methods for tracing the ownership of Bitcoin wallets and the movement of Bitcoins between owners.

7.3 Recommendations for Mixing Services

Centralized mixing services come with the intrinsic risk of reidentification of users by law enforcement seizure of the central coordination server. For this reason, as well as others, there is a clear need for privacy-concerned Bitcoin users to move away from centralized mixing services.

Bitcoin users who wish to take such measures to protect their privacy should make use of decentralized mixing services based on multi-signature transactions, because these services reduce or eliminate the risk of fraud by requiring all users to verify and sign the final transaction before it is valid. users should also insist on services which use an anonymous communications layer to coordinate these transactions, whether that is a cryptographic system such as a DC-net or an anonymous routing system such as Tor.

Academic research in the area of Bitcoin mixing should focus on human factors and usability. Most if not all high-assurance mixing systems currently proposed in the literature present significant practical limitations which are likely to prevent large-scale usage.

7.4 Recommendations for Bitcoin Services

While the applicability of AML regulations to Bitcoin services was questionable very early after the release of Bitcoin, it is now clear that Bitcoin services must meet the full set of AML requirements.

Whenever possible, AML should be performed “at the edge” before traditional currencies are “converted” to Bitcoin (in the language used by regulators). This includes the verification of customer identities and monitoring of their traditional currency transactions for abnormal behavior.

The funds travel rule requires that Bitcoin services forward customer information alongside any funds sent to other Bitcoin services. This rule suggests that there may be significant value in the development of a standardized technical solution for securely exchanging customer identification information alongside Bitcoin transactions. Such a system for customer information exchange would significantly reduce the risk associated with transactions that include identification.

A natural consequence of a Bitcoin identity clearinghouse would be increasing reliance on known customer identities. Bitcoin services may start to refuse to accept transactions which do not come with verified identification information, since they present a much higher risk than those that do. Such a situation could lead to further reduction of the degree of privacy offered by Bitcoin as exchanging identification becomes a constant part of Bitcoin usage.

Bitcoin services should make use of risk scoring technology to monitor Bitcoin transactions for suspicious properties. Bitcoins which appear to have recently been laundered should not be accepted as the customer presents a very high risk of being involved in money laundering. At the same time, service providers must be aware that the Bitcoin privacy community is aware of such transaction risk systems and may game them to reduce their efficacy.

7.5 Recommendations for AML Regulation

Bitcoin AML programs focused on monitoring traditional currency transactions at the edge of the Bitcoin economy are simpler to implement but of limited utility, considering that it is now quite possible for crime to occur entirely within the space of the Bitcoin economy.

It is critical for regulators to develop a clear set of requirements for regulated businesses which operate entirely within Bitcoin. The current system of proprietary, commercial risk-scoring algorithms likely leads to inconsistent and difficult to verify measures taken to prevent laundering. However, Bitcoin services have few opportunities to do better.

AML regulations applied to Bitcoin should require that Bitcoin services maintain a complete CIP and forward identification information with all transactions sent to other Bitcoin services. These regulations must apply equally to exchanges and to Bitcoin services which are less obviously part of the finance system, such as retailers, since they can be used as part of the money laundering process.

The volume and complexity of Bitcoin transactions requires that Bitcoin AML regulation be implemented and enforced by automated means. SARs and other AML reporting should be handled in machine-readable form and FinCEN should invest in algorithmic analysis using data science methods.

Rules developed to prevent money laundering within Bitcoin will need to be strictly enforced on a global scale, as the services and organizations currently involved in Bitcoin laundering appear to have already largely moved to countries with little or no AML enforcement. To date, law enforcement has taken action against Bitcoin services enabling laundering in only a handful of cases, leading to a perception that facilitating laundering is a profitable venture with limited risk.

7.6 Fundamental Conflict

To a large extent, the goals of the Bitcoin community of users are at odds with the goals of AML regulation. This conflict manifests in the political and business realms, as Bitcoin services push back against AML regulation by locating in jurisdictions with limited AML enforcement, designing their services in such a way as to minimize AML obligations, or simply failing to comply with regulation until forced to do so or shut down.

This conflict also manifests on a technical level as an ongoing evolution of anonymization and blockchain analysis technology. Improved methods for Bitcoin mixing continue to be developed, while forensic analysts and risk analysis services develop improved methods of blockchain analysis.

However, several factors discourage large-scale success on the part of Bitcoin mixing services. First, users encounter risk and frustration in attempts to use Bitcoin mixing services, and the level of difficulty involved in use of a mixing service tends to increase as the security and trustworthiness of the system improve, due to the requirement for client-side programs. These factors discourage the prevalent use of mixing techniques, suggesting that mixing will likely remain a minority behavior for the foreseeable future.

Additionally, increasingly effective global AML regulations directed towards Bitcoin services will provide further disincentive for mixing. As more Bitcoin exchanges require customer identification and refuse to accept Bitcoin which has been laundered, it will become difficult to “cash out” or otherwise use laundered Bitcoin. Due to the complexities of international financial regulation this process is slow but has already shown some success with several prominent cases of law enforcement intervention in both exchanges and mixers which failed to take sufficient AML measures.

Ultimately, Bitcoin is being advertised and used for a purpose for which it is not well suited. Bitcoin is not an anonymous system, it is a pseudonymous system at best—and it is broadly understood that pseudonymity is usually easily broken.

Other cryptocurrencies designed with the explicit goal of privacy, such as Monero and Dash, are far more suited to the needs of users who desire complete privacy. On the other hand, these cryptocurrencies present a near worst-case scenario to financial regulators, as AML regulations applied at the edge will become the only opportunity to deter money laundering.

As was the case prior to the passage of the Banking Secrecy Act, tracing of the origin of money may simply not be practical in the future. Law enforcement and regulators will need to focus on other methods of detecting and investigating financial crime.

A. LIST OF MIXING SERVICES

Table A.1 lists all mixing services considered. For Disposition, “current” indicates a currently active service. “SDBO” indicates a service which was shut down by its owner. “NI” indicates a service which failed to garner sufficient interest to be discussed or reviewed, and which is no longer operating. Finally, “scam?” indicates a service which appears to have been a scam, based on user reports.

Table A.1: Mixing services evaluated

Name	Date Announced	Type	Disposition	Darknet
bitcoinfo.com	2011-10-27	central	Scam?	No
easywallet.org	2012-04-06	central	SDBO	No
MixMyCoin	2013-11-21	central	Scam?	No
BitLaunder	2014-01-04	central	Scam?	No
bitmixer.io	2014-01-14	central	SDBO	No
bitcoin blender	2014-01-28	central	Scam?	
coinclean.cc	2014-03-05	central	SDBO	Yes
BitiMix	2014-06-05	central	NI	Yes
bitmix	2014-06-13	central	NI	No
BctMixPot.com	2014-07-28	central	NI	No
btctumbler	2014-08-20	central	NI	Yes
btcmix.me	2014-08-29	central	NI	No
btcmixing.com	2014-09-16	central	NI	Yes
JoinMarket	2015-01-09	coinjoin	NI	Yes
mixing.space	2015-03-04	central	NI	Yes
coinmixer.se	2015-06-13	central	sdbo	Yes
anonymizer5lg2fz.onion	2015-07-23	central	SDBO	Yes
coinmixer.net	2015-08-19	central	Scam?	Yes

Table A.1: (continued)

darklaunder.com	2015-08-19	central	Scam?	No
spacechain.io	2015-08-31	central	Scam?	No
DeepMix	2015-09-09	mixcoin	NI	Exclusively
bitmix.in	2015-09-20	central	Scam?	Yes
Bitcloak	2016-02-21	central	current	Exclusively
bitcloak	2016-02-21	central	current	Yes
cryptomixer.io	2016-03-30	central	current	Yes
mixcoin.tk	2016-11-27	central	Scam?	Yes
mixem.io	2016-12-19	central	NI	No
mixer.money	2016-12-22	central	current	Exclusively
Burger	2017-02-17	central	NI	No
chipmixer	2017-05-26	chip	current	Yes
gocrypto.io	2017-05-29	central	NI	Yes
FRECOINLAUNDRY	2017-07-03	central	NI	Exclusively
5ifblitg2ywjjo2t.onion	2017-07-31	central	NI	Yes
bitmixer.co	2017-08-05	central	SDBO	No
bitmix.biz	2017-08-18	central	current	Yes
bitmix.biz	2017-08-18	central	current	Yes
btcmixer.biz	2017-08-24	central	NI	Yes
privcoin.io	2017-09-09	central	current	No
bitmixcoin.io	2017-09-10	central	Scam?	Yes
mixer.to	2017-09-16	central	Scam?	Yes
foxmixer	2017-10-01	central	current	Yes
http://2ovmq6sfab6u4ucr.onion/	2017-10-26	central	Scam?	No
smartmix.io	2018-01-24	central	current	Yes

Table A.1: (continued)

bestmixer.io	2018-03-16	central	takedown	Yes
bitwhisk.io	2018-03-27	central	SDBO	
bitsafe.pro	2018-05-03	central	Scam?	No
bitmaximum.io	2018-06-30	central	SDBO	Yes
jambler.io	2018-07-13	central	current	Yes
doublemixer.com	2018-08-16	central	Scam?	No
mixtum.io	2018-08-27	central	current	Yes
wasabi wallet	2018-09-24	coinjoin	current	Yes
ecomix.io	2018-09-25	central	SDBO	Yes
blender.io	2018-10-18	central	current	Yes
coinmixer.be	2018-11-20	central	Scam?	Yes
bitcoin laundry	2018-12-03	central	Scam?	Yes
zatoshi	2018-12-05	coinjoin	NI	Yes
domixer	2019-03-15	central	NI	Yes
bitcoin.dj	2019-03-27	central	NI	No
bitmixer.xyz	2019-04-17	central	Scam?	Yes
btcanonmixer.com	2019-04-18	central	Scam?	Yes
mybitmix.com	2019-05-16	central	current	Exclusively
mixertumbler.com	2019-05-21	central	current	Yes
bmc	2019-07-17	central	NI	Exclusively
smartmixer.io	2019-07-24	central	current	Yes
sharedcoin.com		coinjoin	SDBO	Yes
Fogify		central	shut down	Exclusively
pay shield		central	current	
Helix Light		central	Scam?	No

Table A.1: (continued)

Helix Grams	central	SDBO	Yes
-------------	---------	------	-----

B. LIST OF ANALYSIS SERVICES

Comparatively few AML transaction risk scoring services exist, and they generally publish very little about their internal operation. Table B.1 lists services which were evaluated for advertised capabilities and methods.

Table B.1 AML risk scoring services

Name	Webpage
BitRank	https://bitrankverified.com/
Scorechain	https://www.scorechain.com/
Elliptic	https://www.elliptic.co/
CipherTrace	https://ciphertrace.com/
Chainalysis	https://www.chainalysis.com/