A Fistful of Bitcoins: Characterizing Payments Among Men with No Names

[Extended Abstract]

Sarah Meiklejohn[*]
University College London
s.meiklejohn@ucl.ac.uk

Marjori Pomarole
marjoripomarole@gmail.com

Grant Jordan
gejordan@cs.ucsd.edu

Kirill Levchenko
UC San Diego
klevchen@cs.ucsd.edu

Damon McCoy
ICSI
damon.mccoy@gmail.com

Geoffrey M. Voelker
UC San Diego
voelker@cs.ucsd.edu

Stefan Savage
UC San Diego
savage@cs.ucsd.edu

ABSTRACT

Bitcoin is a purely online virtual currency, unbacked by either physical commodities or sovereign obligation; instead, it relies on a combination of cryptographic protection and a peer-to-peer protocol for witnessing settlements. Consequently, Bitcoin has the unintuitive property that while the ownership of money is implicitly anonymous, its flow is globally visible. In this paper we explore this unique characteristic further, using heuristic clustering to group Bitcoin wallets based on evidence of shared authority, and then using re-identification attacks (i.e., empirical purchasing of goods and services) to classify the operators of those clusters. From this analysis, we consider the challenges for those seeking to use Bitcoin for criminal or fraudulent purposes at scale.

1. INTRODUCTION

Demand for low friction e-commerce of various kinds has driven a proliferation in online payment systems over the last decade. Thus, in addition to established payment card networks (e.g., Visa and Mastercard) a broad range of so-called "alternative payments" has emerged including eWallets (e.g., Paypal, Google Checkout, and WebMoney), direct debit systems (typically via ACH, such as eBillMe), money transfer systems (e.g., Moneygram) and so on. However, virtually all of these systems have the property that they are denominated in existing fiat currencies (e.g., dollars), explicitly identify the payer in transactions, and are centrally or quasi-centrally administered. (In particular, there is a central controlling authority who has the technical and legal capacity to tie a transaction back to a pair of individuals.)

By far the most intriguing exception to this rule is Bitcoin. First deployed in 2009, Bitcoin is an independent online monetary system that combines some of the features of cash and existing online payment methods. Like cash, Bitcoin transactions do not explicitly identify the payer or the payee: a transaction is a cryptographically signed transfer of funds from one public key to another. Moreover, like cash, Bitcoin transactions are irreversible (in particular, there is no *chargeback* risk as with credit cards). However, unlike cash, Bitcoin requires third-party mediation: a global peer-to-peer network of participants validates and certifies all transactions. Such decentralized accounting requires each network participant to maintain the entire transaction history of the system, which even in 2012 amounted to over 3GB of compressed data. Bitcoin identities are thus *pseudo-anonymous*: while not explicitly tied to real-world individuals or organizations, all transactions are completely transparent.

This unusual combination of features has given rise to considerable confusion about the nature and consequences of the anonymity that Bitcoin provides. In particular, there is concern that the combination of scalable, irrevocable, anonymous payments would prove highly attractive for criminals engaged in fraud or money laundering. In a widely leaked 2012 Intelligence Assessment, FBI analysts make just this case and conclude that a key "advantage" of Bitcoin for criminals is that "law enforcement faces difficulties detecting suspicious activity, identifying users and obtaining transaction records" [5]. Similarly, in a late 2012 report on Virtual Currency Schemes, the European Central Bank opines that the lack of regulation and due diligence might enable "criminals, terrorists, fraudsters and money laundering" and that "the extent to which any money flows can be traced back to a particular user is unknown" [4]. Indeed, there is at least some anecdotal evidence that this statement is true, with the widely publicized Silk Road service using Bitcoin to trade in a range of illegal goods (e.g., restricted drugs and firearms). Finally, adding to this urgency is Bitcoin's considerable growth, both quantitatively — a merchant servicer, Bitpay, announced that it had signed up over 1,000

The original version of this paper is entitled "A Fistful of Bitcoins: Characterizing Payments Among Men with No Names" and was published in the proceedings of the Internet Measurement Conference, 2013, ACM.

[*]Work done while a graduate student at UC San Diego.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 2008 ACM 0001-0782/08/0X00 ...$5.00.

merchants in 2012 to accept bitcoins, and in November 2013 the exchange rate soared to a peak of 1,000 USD per bitcoin — and qualitatively via integration with existing payment mechanisms and the increasing attention of world financial institutions. In 2012 alone, Bitinstant offered to tie users' Bitcoin wallets to Mastercard accounts [3], Bitcoin Central partnered with the French bank Crédit Mutuel Arkéa to gateway Bitcoin into the banking system [8], Canada decided to tax Bitcoin transactions [2], and FinCEN issued regulations on virtual currencies [6]). Despite this background of intense interest, Bitcoin's pseudo-anonymity has limited how much is known about how the currency is used and how Bitcoin's use has evolved over time.

In this context, our work seeks to better understand the traceability of Bitcoin flows. Importantly, our goal is not to generally de-anonymize all Bitcoin users — as the abstract protocol design itself dictates that this should be impossible — but rather to identify certain *idioms of use* present in concrete Bitcoin network implementations that erode the anonymity of the users who engage in them. We stress that our work was done at a specific point in the evolution of Bitcoin, and that as idioms of use change, the techniques we develop may need to adapt as well.

Our approach is based on the availability of the Bitcoin *block chain*: a replicated graph data structure that encodes all Bitcoin activity, past and present, in terms of the public digital signing keys party to each transaction. However, since each of these keys carries no explicit information about ownership, our analysis depends on imposing additional structure on the transaction graph.

Our methodology has two phases. First, in Section 3, we describe a re-identification attack wherein we open accounts and make purchases from a broad range of known Bitcoin merchants and service providers. Since one endpoint of the transaction is known (i.e., we know which public key we used), we are able to positively label the public key on the other end as belonging to the service; we augment this attack by crawling Bitcoin forums for "self-labeled" public keys (e.g., where an individual or organization explicitly advertises a key as their own). Next, in Section 4, we build on past efforts [1, 9, 10, 12] to cluster public keys based on evidence of shared spending authority. This clustering allows us to amplify the results of our re-identification attack: if we labeled one public key as belonging to a particular service, we can now transitively taint the entire cluster containing this public key as belonging to that service as well. The result is a condensed graph, in which nodes represent entire users and services rather than individual public keys.

From this data, we examine the suitability of Bitcoin for hiding large-scale illicit transactions. Using the dissolution of a large Silk Road wallet and notable Bitcoin thefts as case studies, we argue that an agency with subpoena power would be well placed to identify who is paying money to whom. Indeed, we argue that the increasing dominance of a small number of Bitcoin institutions (most notably services that perform currency exchange), coupled with the public nature of transactions and our ability to label monetary flows to major institutions, ultimately makes Bitcoin unattractive for high-volume illicit use such as money laundering.

2. BITCOIN BACKGROUND

The heuristics that we use to cluster pseudonyms depend on the structure of the Bitcoin protocol, so we first describe it here, and briefly mention the anonymity that it is intended to provide. Additionally, much of our analysis discusses the "major players" and different categories of Bitcoin-based services, so we also present a more high-level overview of Bitcoin participation.

2.1 Bitcoin protocol description

Bitcoin is a decentralized electronic currency, introduced by (the pseudonymous) Satoshi Nakamoto in 2008 [7] and deployed on January 3 2009. Briefly, a bitcoin can be thought of as a chain of *transactions* from one owner to the next, where owners are identified by a *public key* — from here on out, an *address* — that serves as a pseudonym; i.e., users can use any number of addresses and their activity using one set of addresses is not inherently tied to their activity using another set, or to their real-world identity. In each transaction, the previous owner signs — using the secret signing key corresponding to his address — a hash of the transaction in which he received the bitcoins and the address of the next owner. (In fact, transactions can have many input and output addresses, a fact that we exploit in our clustering heuristics in Section 4, but for simplicity we restrict ourselves here to the case of a single input and output.) This signature (i.e., transaction) can then be added to the set of transactions that constitutes the bitcoin; because each of these transactions references the previous transaction (i.e., in sending bitcoins, the current owner must specify where they came from), the transactions form a chain. To verify the validity of a bitcoin, a user can check the validity of each of the signatures in this chain.

To prevent double spending, it is necessary for each user in the system to be aware of all such transactions. Double spending can then be identified when a user attempts to transfer a bitcoin after he has already done so. To determine which transaction came first, transactions are grouped into *blocks*, which serve to timestamp the transactions they contain and vouch for their validity. Blocks are themselves formed into a chain, with each block referencing the previous one (and thus further reinforcing the validity of all previous transactions). This process yields a *block chain*, which is then publicly available to every user within the system.

This process describes how to transfer bitcoins and broadcast transactions to all users of the system. Because Bitcoin is decentralized and there is thus no central authority minting bitcoins, we must also consider how bitcoins are generated in the first place. In fact, this happens in the process of forming a block: each accepted block (i.e., each block incorporated into the block chain) is required to be such that, when all the data inside the block is hashed, the hash begins with a certain number of zeroes. To allow users to find this particular collection of data, blocks contain, in addition to a list of transactions, a *nonce*. (We simplify the description slightly to ease presentation.) Once someone finds a nonce that allows the block to have the correctly formatted hash, the block is then broadcast in the same peer-to-peer manner as transactions. The system is designed to generate only 21 million bitcoins in total. Finding a block currently comes with an attached reward of 25 BTC; this rate was 50 BTC until November 28 2012 (block height 210,000), and is expected to halve again in 2016, and eventually drop to 0 in 2140.

The dissemination of information within the Bitcoin network is summarized in Figure 1.

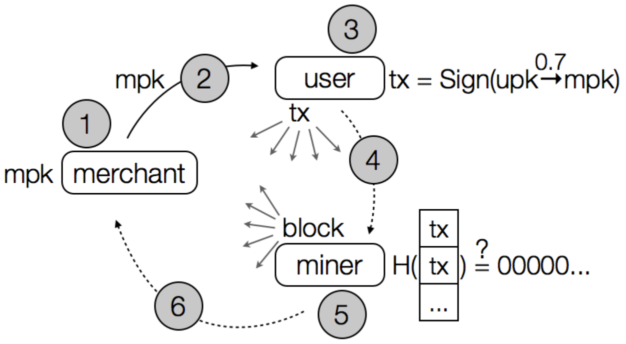

1
2
3
4
5
6
mpk
mpk
merchant
user
tx = Sign(upk $\xrightarrow{0.7}$ mpk)
tx
block
miner
H(tx tx ...) $\overset{?}{=}$ 00000...

Figure 1: How a Bitcoin transaction works; in this example, a user wants to send 0.7 bitcoins as payment to a merchant. In (1), the merchant generates or picks an address *mpk*, and in (2) it sends this address to the user. In (3), the user forms the transaction *tx* to transfer the 0.7 BTC from *upk* to *mpk*. In (4), the user broadcasts this transaction to his peers, which (if the transaction is valid) allows it to flood the network. In this way, a miner learns about his transaction. In (5), the miner works to incorporate this and other transactions into a block by checking if their hash is within some target range. In (6), the miner broadcasts this block to her peers, which (if the block is valid) allows it to flood the network. In this way, the merchant learns that the transaction has been accepted into the global block chain, and thus receives the user's payment.

2.2 Participants in the Bitcoin network

In practice, the way in which Bitcoin can be used is much simpler than the above description might indicate. First, generating a block is so computationally difficult that very few individual users attempt it on their own. Instead, users may join a *mining pool*, in which they contribute "shares" to narrow down the search space, and earn a small amount of bitcoins in exchange for each share.

Users may also avoid coin generation entirely, and simply purchase bitcoins through one of the many *exchanges*. They may then keep the bitcoins in a wallet stored on their computer or, to make matters even easier, use a *wallet service* (although many wallet services have suffered thefts and been shut down).

Finally, to actually spend their bitcoins, users could gamble with one of the popular dice games such as Satoshi Dice. They could also buy items from various online vendors. Finally, users wishing to go beyond basic currency speculation can invest their bitcoins with firms such as Bitcoinica (shut down after a series of thefts) or Bitcoin Savings & Trust (later revealed as a major Ponzi scheme).

3. DATA COLLECTION

To identify addresses belonging to the types of services mentioned in Section 2.2, we sought to "tag" as many addresses as possible; i.e., label an address as being definitively controlled by some known real-world user. As we will see in Section 4.1, by clustering addresses based on evidence of shared control, we can bootstrap off the minimal ground truth data this provides to tag entire clusters of addresses as also belonging to that user.

Our predominant method for tagging users was simply transacting with them (e.g., depositing into and withdrawing bitcoins from Mt. Gox) and then observing the addresses they used. We additionally collected known (or assumed) addresses that we found in various forums and other Web sites, although we regarded this latter kind of tagging as less reliable than our own observed data.

3.1 From our own transactions

We engaged in 344 transactions with a wide variety of services, listed in Table 1, including mining pools, wallet services, bank exchanges, non-bank exchanges, vendors, gambling sites, and miscellaneous services.

Mining pools. We mined bitcoins using an AMD Radeon HD 7970, capable of approximately 530 million SHA-256 computations per second, which allowed us to trigger a payout of at least 0.1 BTC with 11 different pools, anywhere from 1 to 25 times. For each payout transaction, we then labeled the input addresses as belonging to the pool.

Wallets. We kept money with most of the major wallet services (10 in total), and made multiple deposit and withdrawal transactions for each.

Bank exchanges. Most of the real-time trading exchanges (i.e., in which the exchange rate is not fixed) also function as banks. As such, we tagged these services just as we did the wallets: by depositing into and withdrawing from our accounts. We kept accounts with 18 such exchanges in total.

Non-bank exchanges. In contrast, most of the fixed-rate exchanges did not function as banks, and are instead intended for one-time conversions. We therefore were able to participate in fewer transactions with these exchanges, although we again tried to transact with most of the major ones at least once (8 in total).

Vendors. We purchased goods, both physical and digital, from a wide variety of vendors. Many of the vendors we interacted with did not use an independent method for accepting bitcoins, but relied instead on the BitPay payment gateway (and one used WalletBit as a payment gateway). We also kept a wallet with Silk Road, which allowed us to tag their addresses without making any purchases.

Gambling. We kept accounts with five poker sites, and transacted with eight sites offering mini-games and/or lotteries.

Miscellaneous. Four of the additional services we interacted with were *mix* or *laundry* services: when provided with an output address, they promised to send to that address coins that had no association with the ones sent to them; the more sophisticated ones offered to spread the coins out over various transactions and over time. One of these, BitMix, simply stole our money, while Bitcoin Laundry twice sent us our own coins back, indicating we were possibly their only customer at that time. We also interacted with Bit Visitor, a site that paid users to visit certain sites; Bitcoin Advertisers, which provided online advertising; CoinAd, which gave out free bitcoins; Coinapult, which forwarded bitcoins to an email address, where they could then be redeemed; and finally, Wikileaks, with whom we donated to both their

Mining		
50 BTC	BTC Guild	Itzod
ABC Pool	Deepbit	Ozcoin
Bitclockers	EclipseMC	Slush
Bitminter	Eligius	
Wallets		
Bitcoin Faucet	Easywallet	Strongcoin
My Wallet	Flexcoin	WalletBit
Coinbase	Instawallet	
Easycoin	Paytunia	
Exchanges		
Bitcoin 24	BTC-e	Aurum Xchange
Bitcoin Central	CampBX	BitInstant
Bitcoin.de	CA VirtEx	Bitcoin Nordic
Bitcurex	ICBit	BTC Quick
Bitfloor	Mercado Bitcoin	FastCash4Bitcoins
Bitmarket	Mt Gox	Lilion Transfer
Bitme	The Rock	Nanaimo Gold
Bitstamp	Vircurex	OKPay
BTC China	Virwox	
Vendors		
ABU Games	BTC Buy	HealthRX
Bitbrew	BTC Gadgets	JJ Games
Bitdomain	Casascius	NZBs R Us
Bitmit	Coinabul	Silk Road
Bitpay	CoinDL	WalletBit
Bit Usenet	Etsy	Yoku
Gambling		
Bit Elfin	BitZino	Gold Game Land
Bitcoin 24/7	BTC Griffin	Satoshi Dice
Bitcoin Darts	BTC Lucky	Seals with Clubs
Bitcoin Kamikaze	BTC on Tilt	
Bitcoin Minefield	Clone Dice	
Miscellaneous		
Bit Visitor	Bitfog	CoinAd
Bitcoin Advertisers	Bitlaundry	Coinapult
Bitcoin Laundry	BitMix	Wikileaks

Table 1: The various services we interacted with, grouped by (approximate) type.

public donation address and two one-time addresses generated for us via their IRC channel.

3.2 From other sources

In addition to our own transactions, many users publicly claim their own addresses; e.g., charities providing donation addresses, or LulzSec claiming their address on Twitter. While we did not attempt to collect all such instances, many of these tags are conveniently collected at `blockchain.info/tags`, including both addresses provided in users' signatures for Bitcoin forums, as well as self-submitted tags. We collected all of these tags — over 5,000 in total — keeping in mind that the ones that were not self-submitted (and even the ones that were) could be regarded as less reliable than the ones we collected ourselves.

Finally, we searched through the Bitcoin forums (in particular, `bitcointalk.org`) looking for addresses associated with major thefts, or now-defunct services such as Tradehill and GLBSE. Again, these sources are less reliable, so we consequently labeled users only for addresses for which we could gain some confidence through manual due diligence.

4. ADDRESS CLUSTERING

In this section, we present two heuristics for linking addresses controlled by the same user, with the goal of collapsing the many addresses seen in the block chain into larger entities. The first heuristic, in which we treat different addresses used as inputs to a transaction as being controlled by the same user, has already been used and explored in previous work, and exploits an inherent property of the Bitcoin protocol. The second is new and based on so-called *change addresses*; in contrast to the first, it exploits a current *idiom of use* in the Bitcoin network rather than an inherent property. As such, it is less robust in the face of changing patterns within the network, but — as we especially see in Section 5 — it can provide insight into the Bitcoin network that the first heuristic does not.

4.1 Our heuristics

Heuristic 1.

The first heuristic, in which we link together addresses used as input to the same transaction, has already been used many times in previous work [1, 9, 10, 12]. For completeness, we nevertheless present it here as Heuristic 1: if two (or more) addresses are used as inputs to the same transaction, then they are controlled by the same user.

Using this heuristic, we partitioned the network into 5.5 million clusters of users. By naming these clusters — using the data collection described in Section 3 — we observed that some of them corresponded to the same user; e.g., there were 20 clusters that we tagged as being controlled by Mt. Gox. (This is not surprising, as many big services appear to spread their funds across a number of distinct addresses to minimize the risk in case any one gets compromised.) Factoring in "sink" addresses that have to date never sent any bitcoins (and thus did not get clustered using this heuristic) yields at most 6,595,564 distinct users, although we consider this number a quite large upper bound.

Heuristic 2.

Although Heuristic 1 already yields a useful clustering of users, restricting ourselves to only this heuristic does not tell the whole story. To further collapse users, our second heuristic focuses on the role of change addresses within the Bitcoin system. A similar heuristic was explored by Androulaki et al. [1] (who called them "shadow" addresses), although there are a number of important differences. In particular, their definition of shadow addresses relied upon assumptions that may have held at the time of their work, but no longer hold at present. For example, they assumed that users rarely issue transactions to two different users, which is a frequent occurrence today (e.g., payouts from mining pools, or bets on gambling sites).

One of the defining features of the Bitcoin protocol is the way that bitcoins must be spent. When the bitcoins redeemed as the output of a transaction are spent, they must be spent all at once: the only way to divide them is through the use of a *change address*, in which the excess from the input address is sent back to the sender. In one idiom of use, the change address is created internally by the Bitcoin client and never re-used; as such, a user is unlikely to give out this change address to other users (e.g., for accepting payments), and in fact might not even know the address unless he inspects the block chain. If we can identify change addresses, we can therefore potentially cluster not only the input addresses for a transaction (according to Heuristic 1) but also the change address and the input user.

Because our heuristic takes advantage of this idiom of use, rather than an inherent property of the Bitcoin protocol, it does lack robustness in the face of changing (or adversarial) patterns in the network. Furthermore, it has one very negative potential consequence: falsely linking even a small number of change addresses might collapse the entire graph into large "super-clusters" that are not actually controlled by a single user (in fact, we see this exact problem occur in Section 4.2). We therefore focused on designing the safest heuristic possible, even at the expense of losing some utility by having a high false negative rate, and acknowledge that such a heuristic might have to be redesigned or ultimately discarded if habitual uses of the Bitcoin protocol change significantly.

Working off the assumption that a change address has only one input (again, as it is potentially unknown to its owner and is not re-used by the client), we first looked at the outputs of every transaction. If only one of the outputs met this pattern, then we identified that output as the change address. If, however, multiple outputs had only one input and thus the change address was ambiguous, we did not label any change address for that transaction. We also avoided certain transactions; e.g., in a coin generation, none of the outputs are change addresses.

In addition, in custom usages of the Bitcoin protocol it is possible to specify the change address for a given transaction. Thus far, one common usage of this setting that we have observed has been to provide a change address that is in fact the same as the input address. (This usage is quite common: 23% of all transactions in the first half of 2013 used self-change addresses.) We thus avoid such "self-change" transactions as well.

To bring all of these behaviors together, we say that an address is a *one-time change address* for a transaction if the following four conditions are met: (1) the address has not appeared in any previous transaction; (2) the transaction is not a coin generation; (3) there is no self-change address; and (4) all the other output addresses in the transaction have appeared in previous transactions. Heuristic 2 then says that the one-time change address — if one exists — is controlled by the same user as the input addresses.

4.2 Refining Heuristic 2

Although effective, Heuristic 2 is more challenging and significantly less safe than Heuristic 1. In our first attempt, when we used it as defined above, we identified over 4 million change addresses. Due to our concern over its safety, we sought to approximate the false positive rate. To do this even in the absence of significant ground truth data, we used the fact that we could observe the behavior of addresses over time: if an address looked like a one-time change address at one point in time (where time was measured by block height), and then at a later time the address was used again, we considered this a false positive. Stepping through time in this manner allowed us to identify 555,348 false positives, or 13% of all labeled change addresses.

We then considered ways of making the heuristic more conservative. First, however, a manual inspection of some of these false positives revealed an interesting pattern: many of them were associated with transactions to and from Satoshi Dice and other dice games. By looking further into the payout structure of these games, it became clear that these were not truly false positives, as when coins are sent to Satoshi Dice, the payout is sent back to the same address. If a user therefore spent the contents of a one-time change address with Satoshi Dice, the address would receive another input back from Satoshi Dice, which would appear to invalidate the "one-timeness" of the address. We therefore chose to ignore this case, believing that addresses that received later inputs solely from Satoshi Dice could still be one-time change addresses. By doing so the false positive rate reduces to only 1%. We next considered waiting to label an address as a change address; i.e., waiting to see if it received another input. Waiting a day drove the false positive rate down to 0.28%; waiting a week drove it down to 0.17%, or only 7,382 false positives total.

Despite all these precautions, when we clustered users using this modified heuristic, we still ended up with a giant super-cluster containing the addresses of Mt. Gox, Instawallet, BitPay, and Silk Road, among others; in total, this super-cluster contained 1.6 million addresses. After a manual inspection of some of the links that led to this super-cluster, we discovered two problematic patterns. First, especially within a short window of time, the same change address was sometimes used twice. Second, certain addresses were occasionally used as "self-change" addresses, and then later used as separate change addresses. We thus further refined our heuristic by ignoring transactions involved with either of these types of behavior. For transactions in which an output address had already received only one input, or for transactions in which an output address had been previously used in a self-change transaction, we chose to not tag anything as the change address. Doing so, and manually removing a handful of other false positives (with no discernible pattern), we identified 3,540,831 change addresses.

Using this refined Heuristic 2 produces 3,384,179 clusters, which we were able to again collapse slightly (using our tags) to 3,383,904 distinct clusters. Of these clusters, we were able to name 2,197 of them (accounting for over 1.8 million addresses). Although this might seem like a small fraction, recall that by participating in 344 transactions we hand-tagged only 1,070 addresses, and thus Heuristic 2 allowed us to name 1,600 times more addresses than our own manual observation provided. Furthermore, as we will argue in Section 5, the users we were able to name capture an important and active slice of the Bitcoin network.

Having finally convinced ourselves of the safety of Heuristic 2, by refining it substantially, and its effectiveness, we use Heuristic 2 exclusively for the results in the next section.

5. ANALYSIS OF ILLICIT ACTIVITY

Exchanges have essentially become chokepoints in the Bitcoin economy, in the sense that it is unavoidable to buy into or cash out of Bitcoin at scale without using an exchange. While sites like `localbitcoins.com` and `bitcoinary.com` do allow users to avoid exchanges (for the former, by pairing buyers directly with sellers in their geographic area), the current and historical volume on these sites does not seem to be high enough to support cashing out at scale.

In this section, we argue that this centrality presents a unique problem for criminals: if a thief steals thousands of bitcoins, this theft is unavoidably visible within the Bitcoin network, and thus the initial address of the thief is known and (as most exchanges try to maintain some air of reputability) he cannot simply transfer the bitcoins directly from the theft to a known exchange. While he might attempt to use a mix service to hide the source of the money, we again argue that these services do not currently have the volume to launder thousands of bitcoins. As such, we explore in this section various alternative strategies that thieves have developed for hiding the source of stolen bitcoins. In particular, we focus on the effectiveness of Heuristic 2 in de-anonymizing these flows, and thus in tracking illicitly obtained bitcoins to exchanges (and thus, e.g., providing an agency with subpoena power the opportunity to learn whose account was deposited into, and in turn potentially the identity of the thief). For this to work, we do not need to (and cannot) account for each and every stolen bitcoin, but rather need to demonstrate only some flow of bitcoins directly from the theft to an exchange or other known institution.

To demonstrate the effectiveness of Heuristic 2 in this endeavor, we focus on an idiom of use that we call a "peeling chain." The usage of this pattern extends well beyond criminal activity, and is seen (for example) in the withdrawals for many banks and exchanges, as well as in the payouts for some of the larger mining pools. In a peeling chain, a single address begins with a relatively large amount of bitcoins (e.g., for mining pools it starts with the 25 BTC reward). A smaller amount is then "peeled" off this larger amount, creating a transaction in which a small amount is sent to one address and the remainder is sent to a one-time change address. This process is repeated — potentially for hundreds or thousands of hops — until the larger amount is pared down. By using Heuristic 2, we are able to track flows of money by following these change links systematically: at each hop, we look at the two output addresses in the transaction. If one of these output addresses is a change address, we can follow the chain to the next hop by following the change address (i.e., the next hop is the transaction in which this change address spends its bitcoins), and can identify the meaningful recipient in the transaction as the other output address (the "peel").

Silk Road and Bitcoin Savings & Trust.

One of the most well-known and heavily scrutinized addresses in Bitcoin's history is`1DkyBEKt` — full address: `1DkyBEKt5S2GDtv7aQw6rQepAvnsRyHoYM` — which is believed to be associated with Silk Road and was active between January and September 2012. Starting in January, the address began to receive large aggregate sums of bitcoins; in the first of these, the funds of 128 addresses were combined to deposit 10,000 BTC into the `1DkyBEKt` address, and many transactions of this type followed. All together, the address received 613,326 BTC in a period of eight months, receiving its last aggregate deposit on August 16 2012.

Then, starting in August 2012, bitcoins were aggregated and withdrawn from `1DkyBEKt`: first, amounts of 20,000, 19,000, and 60,000 BTC were sent to separate addresses; later, 100,000 BTC each was sent to two distinct addresses, 150,000 BTC to a third, and 158,336 BTC to a fourth, effectively emptying the `1DkyBEKt` address of all of its funds.

Due to its large balance (at its height, it contained 5% of all generated bitcoins), as well as the curious nature of its rapidly accumulated wealth and later dissolution, this address has naturally been the subject of heavy scrutiny by the Bitcoin community. While it is largely agreed that the address is associated with Silk Road (and indeed our clustering heuristic did tag this address as being controlled by Silk Road), some have theorized that it was the "hot" (i.e., active) wallet for Silk Road, and that its dissipation represents a changing storage structure for the service. Others, meanwhile, have argued that it was the address belonging to the user pirate@40, who was responsible for carrying out the largest Ponzi scheme in Bitcoin history (the investment scheme Bitcoin Savings & Trust, which is now the subject of a lawsuit brought by the SEC [11]).

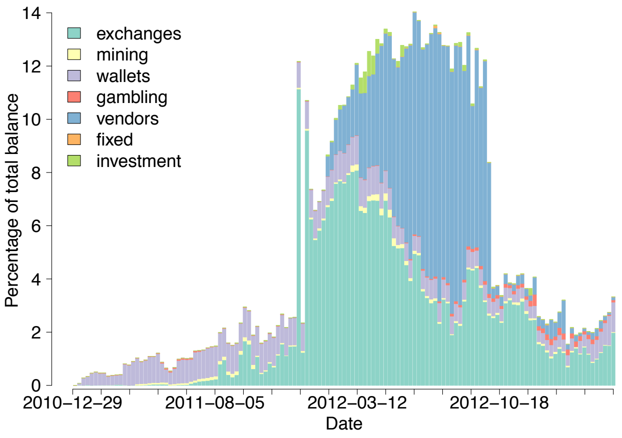

Figure 2: The balance of each major category, represented as a percentage of total active bitcoins; i.e., the bitcoins that are not held in sink addresses.

To see where the funds from this address went, and if they ended up with any known services, we first plotted the balance of each of the major categories of services, as seen in Figure 2. Looking at this figure, it is clear that when the address was dissipated, the resulting funds were not sent en masse to any major services, as the balances of the other categories do not change significantly. To nevertheless attempt to find out where the funds did go, we turn to the traffic analysis described above.

In particular, we focus on the last activity of the `1DkyBEKt` address, when it deposited 158,336 BTC into a single address. This address then peeled off 50,000 BTC each to two separate addresses, leaving 58,336 BTC for a third address; each of these addresses then began a peeling chain, which we followed using the methodology described above (i.e., at each hop we continued along the chain by following the change address, and considered the other output address to be a meaningful recipient of the money). After following 100 hops along each chain, we observed peels to the services listed in Table 2.

In this table, we see that, although a longitudinal look at the balances of major services did not reveal where the money went, following these chains revealed that bitcoins were in fact sent to a variety of services. The overall balance was not highly affected, however, as the amounts sent were relatively small and spread out over a handful of transactions. Furthermore, while our analysis does not itself reveal the owner of `1DkyBEKt`, the flow of bitcoins from this address to known services demonstrates the prevalence of these services (54 out of 300 peels went to exchanges alone) and provides the potential for further de-anonymization: the evidence that the deposited bitcoins were the direct result

	First		Second		Third	
Service	Peels	BTC	Peels	BTC	Peels	BTC
Bitcoin-24			1	2	3	124
Bitcoin Central					2	2
Bitcoin.de					1	4
Bitmarket					1	1
Bitstamp			5	97	1	1
BTC-e					1	250
CA VirtEx	1	3	1	10	3	22
Mercado Bitcoin					1	9
Mt. Gox	11	492	14	70	5	35
OKPay	2	151			1	125
Instawallet	7	39	5	135	2	43
WalletBit	1	1				
Bitzino					2	1
Seals with Clubs	1	8				
Coinabul			1	29		
Medsforbitcoin	3	10				
Silk Road	4	28			5	102

Table 2: Tracking bitcoins from `1DkyBEKt`. Along the first 100 hops of the first, second, and third peeling chains resulting from the withdrawal of 158,336 BTC, we consider the number of peels seen to each service, as well as the total number of bitcoins (rounded to the nearest integer value) sent in these peels. The services are separated into the categories of exchanges, wallets, gambling, and vendors.

of either a Ponzi scheme or the sale of drugs might motivate Mt. Gox or any exchange (e.g., in response to a subpoena) to reveal the account owner corresponding to the deposit address in the peel, and thus provide information to link the address to a real-world user.

Tracking thefts.

To ensure that our analysis could be applied more generally, we turned finally to a broader class of criminal activity in the Bitcoin network: thefts. Thefts are in fact quite common within Bitcoin: almost every major service has been hacked and had bitcoins (or, in the case of exchanges, other currencies) stolen, and some have shut down as a result.

To begin, we used a list of major Bitcoin thefts found at `https://bitcointalk.org/index.php?topic=83794`. Some of the thefts did not have public transactions (i.e., ones we could identify and study in the block chain), so we limited our attention to the ones that did. For each theft, we first found the specific set of transactions that represented the theft; i.e., the set of transactions in which the sender was the service and the recipient was the thief. Starting with these transactions, we did a preliminary manual inspection of the transactions that followed to determine their approximate type: we considered aggregations, in which bitcoins were moved from several addresses into a single one; folding, in which some of the aggregated addresses were not clearly associated with the theft; splits, in which a large amount of bitcoins was split among two or more addresses; and finally peeling chains, in which smaller amounts were peeled off from a succession of one-time change addresses. Our results are summarized in Table 3.

Theft	BTC	Date	Movement	Exchanges?
MyBitcoin	4019	Jun 2011	A/P/S	Yes
Linode	46,648	Mar 2012	A/P/F	Yes
Betcoin	3171	Mar 2012	F/A/P	Yes
Bitcoinica	18,547	May 2012	P/A	Yes
Bitcoinica	40,000	Jul 2012	P/A/S	Yes
Bitfloor	24,078	Sep 2012	P/A/P	Yes
Trojan	3257	Oct 2012	F/A	No

Table 3: Tracking thefts. For each theft, we list (approximately) how many bitcoins were stolen, when the theft occurred, how the money moved after it was stolen, and whether we saw any bitcoins sent to known exchanges. For the movement, we use A to mean aggregation, P to mean a peeling chain, S to mean a split, and F to mean folding, and list the various movements in the order they occurred.

Briefly, the movement of the stolen money ranged from quite sophisticated layering and mixing to simple and easy to follow. Examining thefts therefore provides another demonstration of the potential for anonymity provided by Bitcoin, and the ways in which current usage falls short of this potential. For the thieves who used the more complex strategies, we saw little opportunity to track the flow of bitcoins (or at least do so with any confidence that ownership was staying the same), but for the thieves that did not there seemed to be ample opportunity to track the stolen money directly to an exchange.

One of the easiest thefts to track was from Betcoin, an early gambling site that was shut down after its server was hacked on April 11 2012 and 3,171 BTC were stolen. The stolen bitcoins then sat in the thief's address until March 15 2013 (when the bitcoin exchange rate began soaring), when they were aggregated with other small addresses into one large address that then began a peeling chain. After 10 hops, we saw a peel go to Bitcoin-24, and in another 10 hops we saw a peel go to Mt. Gox; in total, we saw 374.49 BTC go to known exchanges, all directly off the main peeling chain, which originated directly from the addresses known to belong to the thief.

In contrast, some of the other thieves used more sophisticated strategies to attempt to hide the flow of money; e.g., for the Bitfloor theft, we observed that large peels off several initial peeling chains were then aggregated, and the peeling process was repeated. Nevertheless, by manually following this peel-and-aggregate process to the point that the later peeling chains began, we systematically followed these later chains and again observed peels to multiple known exchanges: the third peel off one such chain was 191.09 BTC to Mt. Gox, and in total we saw 661.12 BTC sent to three popular exchanges (Mt. Gox, BTC-e, and Bitstamp).

Even the thief we had the most difficulty tracking, who stole bitcoins by installing a trojan on the computers of individual users, seemed to realize the difficulty of cashing out at scale. Although we were unable to confidently track the flow of the stolen money that moved, most of the stolen money did not in fact move at all: of the 3,257 BTC stolen to date, 2,857 BTC was still sitting in the thief's address, and has been since November 2012.

With these thefts, our ability to track the stolen money provides evidence that even the most motivated Bitcoin users (i.e., criminals) are engaging in idioms of use that allow us to erode their anonymity. While one might argue that thieves could easily thwart our analysis, our observation is that — at least at the time we performed our analysis — none of the criminals we studied seem to have taken such precautions. We further argue that the fairly direct flow of bitcoins from the point of theft to the deposit with an exchange provides some evidence that using exchanges to cash out at scale is inevitable, as otherwise thieves presumably would have avoided this less anonymous method of cashing out. Thus, Bitcoin does not — again, at the time we performed our analysis — seem to provide a particularly easy or effective way to transact large volumes of illicitly obtained money.

6. CONCLUSIONS

In this study, we presented a longitudinal characterization of the Bitcoin network, focusing on the growing gap — due to certain idioms of use — between the potential anonymity available in the Bitcoin protocol design and the actual anonymity that is currently achieved by users. To accomplish this task, we developed a new clustering heuristic based on change addresses, allowing us to cluster addresses belonging to the same user. Then, using a small number of transactions labeled through our own empirical interactions with various services, we identify major institutions. Even our relatively small experiment demonstrates that this approach can shed considerable light on the structure of the Bitcoin economy, how it is used, and those organizations who are party to it.

Although our work examines the current gap between actual and potential anonymity, one might naturally wonder — given that our new clustering heuristic is not fully robust in the face of changing behavior — how this gap will evolve over time, and what users can do to achieve stronger anonymity guarantees. We posit that to completely thwart our heuristics would require a significant effort on the part of the user, and that this loss of usability is unlikely to appeal to all but the most motivated users (such as criminals). Nevertheless, we leave a quantitative analysis of this hypothesis as an interesting open problem.

7. REFERENCES

[1] E. Androulaki, G. Karame, M. Roeschlin, T. Scherer, and S. Capkun. Evaluating User Privacy in Bitcoin. In *Proceedings of Financial Cryptography 2013*, 2013.

[2] CBC News. Revenue Canada says BitCoins aren't tax exempt, Apr. 2013. www.cbc.ca/news/canada/story/2013/04/26/business-bitcoin-tax.html.

[3] B. P. Eha. Get ready for a Bitcoin debit card. CNNMoney, Apr. 2012. money.cnn.com/2012/08/22/technology/startups/bitcoin-debit-card/index.html.

[4] European Central Bank. Virtual Currency Schemes. ECB Report, Oct. 2012. www.ecb.europa.eu/pub/pdf/other/virtualcurrencyschemes201210en.pdf.

[5] Federal Bureau of Investigation. (U) Bitcoin Virtual Currency Unique Features Present Distinct Challenges for Deterring Illicit Activity. Intelligence Assessment, Cyber Intelligence and Criminal Intelligence Section, Apr. 2012. cryptome.org/2012/05/fbi-bitcoin.pdf.

[6] FinCEN. Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies, Mar. 2013. www.fincen.gov/statutes_regs/guidance/pdf/FIN-2013-G001.pdf.

[7] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. bitcoin.org/bitcoin.pdf.

[8] M. Peck. Bitcoin-Central is Now The World's First Bitcoin Bank...Kind Of. IEEE Spectrum: Tech Talk, Dec. 2012. spectrum.ieee.org/tech-talk/telecom/internet/bitcoincentral-is-now-the-worlds-first-bitcoin-bankkind

[9] F. Reid and M. Harrigan. An Analysis of Anonymity in the Bitcoin System. In *Security and Privacy in Social Networks*, pages 197–223. Springer New York, 2013.

[10] D. Ron and A. Shamir. Quantitative Analysis of the Full Bitcoin Transaction Graph. In *Proceedings of Financial Cryptography 2013*, 2013.

[11] Securities and Exchange Commission. SEC Charges Texas Man With Running Bitcoin-Denominated Ponzi Scheme, July 2013. www.sec.gov/News/PressRelease/Detail/PressRelease/1370539730583.

[12] znort987. blockparser. github.com/znort987/blockparser.