

Tracing Transactions Across Cryptocurrency Ledgers

Haaroon Yousaf, George Kappos, and Sarah Meiklejohn

University College London

{h.yousaf, g.kappos, s.meiklejohn}@ucl.ac.uk

Abstract—One of the defining features of a cryptocurrency is that its ledger, containing all transactions that have ever taken place, is globally visible. As one consequence of this degree of transparency, a long line of recent research has demonstrated that—even in cryptocurrencies that are specifically designed to improve anonymity—it is often possible to track flows of money as it changes hands, and in some cases to de-anonymize users entirely. With the recent proliferation of alternative cryptocurrencies, however, it becomes relevant to ask not only whether or not money can be traced as it moves within the ledger of a single cryptocurrency, but if it can in fact be traced as it moves *across* ledgers. This is especially pertinent given the rise in popularity of automated trading platforms such as ShapeShift, which make it effortless to carry out such cross-currency trades. In this paper, we use data scraped from ShapeShift over a six-month period and the data from eight different blockchains in order to explore this question. Beyond developing new heuristics and demonstrating the ability to create new types of links across cryptocurrency ledgers, we also identify various patterns of cross-currency trades and of the general usage of these platforms, with the ultimate goal of understanding whether they serve either a criminal or a profit-driven agenda.

I. INTRODUCTION

For the past decade, cryptocurrencies such as Bitcoin have been touted for their transformative potential, both as a new form of electronic cash and as a platform to “re-decentralize” aspects of the Internet and computing in general. In terms of their role as cash, however, it has been well established by now that the usage of pseudonyms in Bitcoin does not achieve meaningful levels of anonymity [15], [16], [1], [9], [18], which casts doubt on its role as a payment mechanism and had caused some to refer to it as “Twitter for your bank account” [13]. What’s more, the ability to track flows of coins is not limited to Bitcoin; it in fact extends even to so-called “privacy coins” like Dash [8], [10], Monero [11], [6], and Zcash [14], [5] that incorporate features explicitly designed to improve on Bitcoin’s anonymity guarantees.

Traditionally, criminals attempting to cash out their illicit funds would have to move them into exchanges; indeed, most tracking techniques rely on identifying the addresses associated with these exchanges as a way to observe when these deposits happen [9]. More and more frequently, however, these exchanges are implementing KYC/AML policies in order to comply with regulatory requirements, meaning criminals risk revealing their real identities when using them. As an alternative, there have emerged in the past few years frictionless trading platforms such as ShapeShift¹ in which users can trade between cryptocurrencies without having to—for now at least—provide any meaningful form of identification (or in

some cases without creating an account at all).² Changelly³ is a similar platform that requires its users to be verified.

Part of the reason for these trading platforms to exist is the sheer rise in the number of different cryptocurrencies: according to the popular cryptocurrency data tracker CoinMarketCap there were 36 cryptocurrencies in September 2013, only 7 of which had a stated market capitalization of over 1 million USD,⁴ whereas in September 2018 there were 2003 cryptocurrencies, 867 of which had a market capitalization of over 1 million USD. Given this proliferation of new cryptocurrencies and platforms that make it easy to transact across them, it becomes important to consider not just whether or not flows of coins can be tracked within the transaction ledger of a given currency, but also if they can be tracked as coins move across their respective ledgers as well. This is especially important given that there are documented cases of criminals attempting to use these cross-currency trades to obscure the flow of their coins: the WannaCry ransomware operators, for example, were observed using ShapeShift to convert their ransomed bitcoins into Monero [3]. More generally, these services have the potential to offer an insight into the broader cryptocurrency ecosystem and the thousands of currencies it now contains.

In this paper, we initiate an exploration of the usage of these emerging cross-currency trading platforms, and in particular of the potential that they offer in terms of the ability to track flows of coins as they move across different transaction ledgers. Here we rely on three distinct sources of data: the cryptocurrency blockchains themselves, the data collected via our own interactions with these trading platforms, and—as we describe in Section IV—the information offered by the platforms themselves via their public APIs.

We begin in Section V with a categorization of these trading platforms based on the *clusters* of addresses they comprise in each of the blockchains, and on the interactions they have with other types of services such as traditional exchanges. While this already provides a useful macro-level view of their activity, it is not specific to individual usages of the platform and thus offers only a very limited ability to correlate transactions in different cryptocurrency ledgers. To this end, we move on in Section VI to identify the specific on-chain transactions associated with an advertised ShapeShift transaction, which we are able to do with a relatively high degree of success (identifying, on average, the deposit transaction for the input currency 75% of the time and the withdrawal transaction for the output currency 37% of the time). We then describe in Section VII the different transactional patterns that

¹<https://shapeshift.io>

²ShapeShift announced in September 2018 that it would soon allow only users with an account to trade, but this is currently still optional [19].

³<https://changelly.com>

⁴<https://coinmarketcap.com/historical/20130721/>

can be traced by identifying the relevant on-chain transactions, before bringing all the analysis together by applying it to several case studies in Section VIII. Our particular focus in this last section is on the extent to which usage of the ShapeShift platform seems to be motivated by concerns for anonymity, as opposed to just day trading or other profit-driven activity.

II. RELATED WORK

We are not aware of any other research exploring these cross-currency trading platforms, but consider as related all research that explores the level of anonymity achieved by distinct cryptocurrencies. This work is complementary to our own, as the techniques it develops can be combined with ours to track the entire flow of cryptocurrencies as they move both within and across different ledgers.

Much of the earlier research in this vein focused on Bitcoin [15], [16], [1], [9], [18], and operates by adopting the so-called “multi-input” heuristic, which says that all input addresses in a transaction belong to the same entity (be it an individual or a service such as an exchange). While the safety of this heuristic has been somewhat eroded by privacy-enhancing techniques like CoinJoin [7], new techniques have been developed to avoid such false positives [10], and as such it has now been accepted as standard and incorporated into many tools for Bitcoin blockchain analytics, such as Chainalysis⁵ and Elliptic.⁶ Once addresses are clustered together in this manner, the entity can then further be identified using hand-collected tags that form a ground-truth dataset. We adopt both of these techniques in order to analyze the clusters formed by ShapeShift and Changelly in a variety of cryptocurrency blockchains, as described in Section V.

In response to the rise of newer “privacy coins”, a recent line of research has also worked to demonstrate that the deployed versions of these cryptocurrencies have various properties that diminish the level of anonymity they achieve in practice. This includes work targeting Dash [10], [8], Monero [11], [6], and Zcash [14], [5].

In terms of Dash, its main privacy feature is similar to CoinJoin, in which different senders join forces to create a single transaction representing their transfer to a diverse set of recipients. Despite the intention for this to hide which recipient addresses belong to which senders, research has demonstrated that such links can in fact be created based on the value being transacted [10], [8]. Monero, which allows senders to hide which input belongs to them by using “mix-ins” consisting of the keys of other users, is vulnerable to de-anonymization attacks exploiting the (now-obsolete) case in which some users chose not to use mix-ins, or exploiting inferences about the age of the coins used as mix-ins [11], [6]. Finally, Zcash is similar to Bitcoin, but with the addition of a privacy feature called the shielded pool, which can be used to hide the values and addresses of the senders and recipients involved in a transaction. Recent research has shown that it is possible to significantly reduce the anonymity set provided by the shielded pool, by developing simple heuristics for identifying links between hidden and partly obscured transactions [14], [5].

⁵<https://www.chainalysis.com/>

⁶<https://www.elliptic.co/>

III. BACKGROUND

A. Cryptocurrencies

The first decentralized cryptocurrency, Bitcoin, was created by Satoshi Nakamoto in 2008 [12] and deployed in January 2009. At the most basic level, bitcoins are digital assets that can be traded between sets of users without the need for any trusted intermediary. Bitcoins can be thought of as being stored in a public key, which is controlled by the entity in possession of the associated private key. A single user can store their assets across many public keys, which act as pseudonyms with no inherent link to the user’s identity. In order to spend them, a user can form and cryptographically sign a transaction that acts to send the bitcoins to a recipient of their choice. Beyond Bitcoin, other platforms now offer more robust functionality. In particular, Ethereum allows users to deploy *smart contracts* onto the blockchain, which act as stateful programs. These programs can be triggered by transactions, which act to autonomously execute the program on a given set of inputs. The only limitation for such programs and the transactions that trigger them is their complexity: every operation that they perform comes with some associated cost (measured in *gas*, a subcurrency of ether), and there is a maximum amount of gas that may be spent by a single transaction.

B. Digital asset trading platforms

In contrast to a traditional exchange, a digital asset trading platform allows users to move between different cryptocurrencies without needing to set up an account, and thus without needing to follow KYC/AML regulations. Instead, a user approaches the service and selects a supported input currency $curIn$ (i.e., the currency from which they would like to move money) and a supported output currency $curOut$ (the currency which they would like to obtain). A user additionally specifies a destination address $addr_u$ in the $curOut$ blockchain, which is the address to which the output currency will be sent. The service then presents the user with an exchange rate $rate$ and an address $addr_s$ in the $curIn$ blockchain to which to send money, as well as a miner fee fee that accounts for the transaction they must form in the $curOut$ blockchain. The user then sends to this address the amount amt in $curIn$ they wish to convert, and after some delay the service sends the appropriate amount of the output currency to the specified destination address. This means that an interaction with either of these services results in two transactions: one on the $curIn$ blockchain sending amt to $addr_s$, and one on the $curOut$ blockchain sending (roughly) $rate \cdot amt - fee$ to $addr_u$.

This describes an interaction with an abstracted platform. Today, the two best-known examples are ShapeShift and Changelly, although Changelly does require account creation (and ShapeShift plans to soon). Each platform supports dozens of cryptocurrencies, ranging from better-known ones such as Bitcoin and Ethereum to lesser-known ones such as FirstBlood and Clams. Many of the supported cryptocurrencies actually operate as ERC20 or BTC tokens, meaning they run as contracts on top of the Ethereum and Bitcoin blockchains, respectively, rather than as their own standalone platforms. In Section IV, we describe in more depth the operations of these concrete platforms and our own interactions with them.

IV. DATA COLLECTION AND STATISTICS

In this section, we describe our data sources, as well as some preliminary statistics about the data we collected. We begin in Section IV-A by describing our own interactions with Changelly, which is a trading platform with a limited personal API. We then describe in Section IV-B both our own interactions with ShapeShift, and the data we were able to scrape from their public API, which provided us with significant insight into their overall set of transactions. Finally, we describe in Section IV-C our collection of the data backing eight different cryptocurrencies.

A. Changelly

Changelly offers a simple API⁷ that allows registered users to carry out transactions with the service. Using this API, we engaged in 22 transactions, using the most popular ShapeShift currencies (Table I) to guide our choices for curIn and curOut.

While doing these transactions, we observed that they would sometimes take up to an hour to complete. This is because Changelly is not willing to take any double-spending risk, meaning they require users to wait for a set number of confirmations (shown to the user at the time of their transaction) in the curIn blockchain before executing the transfer on the curOut blockchain. We used this observation to guide our choice of parameters in our identification of on-chain transactions in Section VI.

B. ShapeShift

ShapeShift’s API⁸ allows users to execute their own transactions, of which we did 18 in total in order to gain ground-truth data about the internal operation of the service. As with Changelly, we were able to gain some valuable insights about the operation of the platform via these personal interactions. Whereas ShapeShift did not visually present the number of confirmations they waited before releasing the coins on the curOut blockchain, we again observed long delays, which indicate that they were in fact waiting for a sufficient number of confirmations on the curIn blockchain.

Beyond these personal interactions, however, the API provides information on the operation of the service as a whole. Most notably, it provides two separate pieces of information: a list of up to 50 of the most recent transactions that have taken place (for any user), and the current trading rates between any pair of cryptocurrencies.

For the latter, ShapeShift provides the following information for all cryptocurrency pairs (curIn, curOut): the rate, the limit (i.e., the maximum that can be exchanged), the minimum that can be exchanged, and the miner fee (denominated in curOut). For the former, ShapeShift provides information about the 50 most recent transactions of the form: (curIn, curOut, amt, t , id), where the first three of these are as discussed in Section III-B, t is a UNIX timestamp, and id is an internal identifier for this transaction. Notably, ShapeShift does not offer information on the blockchain transactions associated with this exchange, but we discuss in Section VI heuristics to identify this information anyway.

Using a simple Web scraper, we downloaded the transactions and rates every five seconds for close to seven months: from November 27 2017 until June 19 2018. This resulted in a set of 2,254,632 distinct transactions. Interestingly, we noticed that several earlier test transactions we did with the platform did not show up in their list of recent transactions, which suggests that their published transactions may in fact underestimate their overall activity.

1) *ShapeShift profits*: As with other trading platforms (and indeed most exchanges in general), Shapeshift’s profits come from the fees that cryptocurrencies pay in order to be added to their platform, and from any difference between the “true” exchange rate between a pair of cryptocurrencies and the rate that ShapeShift offers. In particular, for every rate $rate_{in-out}$ that Shapeshift offers, they have the potential to profit from it if they could get a better rate elsewhere. While the first type of profit is opaque to us, we can approximate the second type using their advertised rates and data scraped from the cryptocurrency data tracker CoinGecko,⁹ which we used because it provided historical daily USD prices for every cryptocurrency we observed being used on ShapeShift (whereas most other data trackers lacked this historical data for some of the lower-ranked coins). For a given individual transaction tx moving v units of curIn into curOut, we then estimated the profit ShapeShift made on it as

$$p(\text{tx}) = \left(\frac{rate_{in-USD}}{rate_{out-USD}} - rate_{in-out} \right) \cdot v,$$

where $rate_{in-USD}$ and $rate_{out-USD}$ are the daily average rates for curIn and curOut respectively on the day in which the ShapeShift transaction was performed (as scraped from CoinGecko). We then approximated their total profits throughout the months we scraped their data as

$$\sum_{i=1}^n p(\text{tx}_i),$$

where n was the total number of transactions we scraped. After plugging our scraped values into the formulas above, we concluded that the approximate profit ShapeShift made from the difference in rates across the seven months alone was over 200 million USD.

2) *ShapeShift currencies*: In terms of the different cryptocurrencies used in ShapeShift transactions, their popularity was distributed as seen in Figure 1. As this figure depicts, the overall activity of ShapeShift is (perhaps unsurprisingly) correlated with the price of Bitcoin in the same time period.

ShapeShift supports dozens of cryptocurrencies, and in our data we observed the use of 57 different ones. The most commonly used coins were as shown in Table I and in Figure 2.

It is clear that Bitcoin and Ethereum are the most heavily used currencies, which is perhaps not surprising given the relative ease with which they can be exchanged with fiat currencies on more traditional exchanges, and their rank in terms of market capitalization. While it may initially be surprising that Ethereum is so much more popular than Bitcoin, given Bitcoin’s higher price, this might be explained by Ethereum’s role as a general-purpose platform used to support many tokens (rather than a single specific asset).

⁷<https://api-docs.changelly.com/>

⁸<https://info.shapeshift.io/api>

⁹<https://www.coingecko.com/>

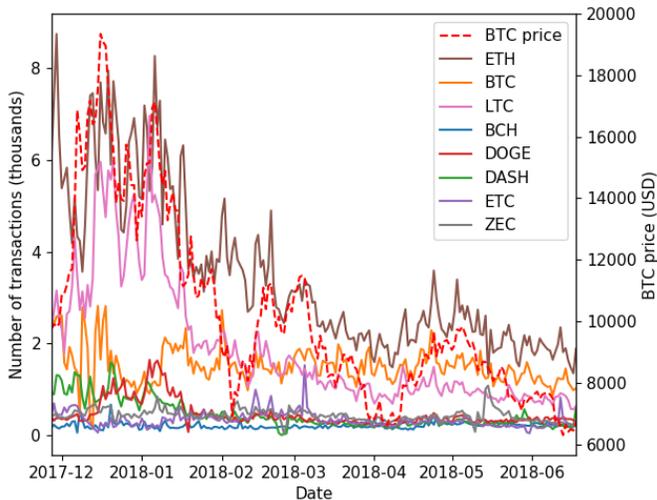


Fig. 1: The total number of transactions per day reported via ShapeShift’s API, and the numbers broken down by cryptocurrency (where a transaction is attributed to a coin if it is used as either curln or curOut). The dotted red line indicates the BTC-USD exchange rate.

Currency	Abbr.	Total	curln	curOut
Ethereum	ETH	1,098,885	710,223	388,662
Bitcoin	BTC	892,147	308,269	583,878
Litecoin	LTC	606,139	390,388	215,751
Bitcoin Cash	BCH	210,687	42,023	168,664
Dogecoin	DOGE	192,615	92,860	99,755
Dash	DASH	151,724	91,015	60,709
Ethereum Classic	ETC	125,739	67,529	58,210
EOS	EOS	125,423	49,594	75,829
Zcash	ZEC	120,386	87,906	32,480
Ripple	XRP	119,363	42,315	77,048

TABLE I: The ten most popular coins used on ShapeShift, in terms of the total units traded, and the respective units traded with that coin as curln and curOut.

C. Blockchain data

For the cryptocurrencies we are interested in exploring further, it is also necessary to download and parse the respective blockchains, in order to identify the transactional behavior of ShapeShift and Changelly. We decided that it was infeasible to do this for all 57 currencies used on ShapeShift (not to mention that given the low volume of transactions, it would likely not yield additional insights anyway), and chose to focus instead on just the top 10, as seen in Table I. Even within the top 10, there were two we chose not to pursue further: Ripple, due to the size of its blockchain (1.7 TB), and EOS, due to the immaturity of its blockchain (it was launched on June 14 2018, and prior to that operated as an Ethereum-based token). This left us with eight cryptocurrencies, which account for 56% of the total number of ShapeShift transactions and 59.3% of the USD value carried by those transactions if we require both curln and curOut to be one of the eight, and 95% of the total number of transactions and 96.6% of the USD value if we require only one of curln or curOut to be one of the eight. Thus, for the remainder of the paper we consider only these eight: Bitcoin, Bitcoin Cash, Dash, Dogecoin, Ethereum, Ethereum Classic, Litecoin, and Zcash.

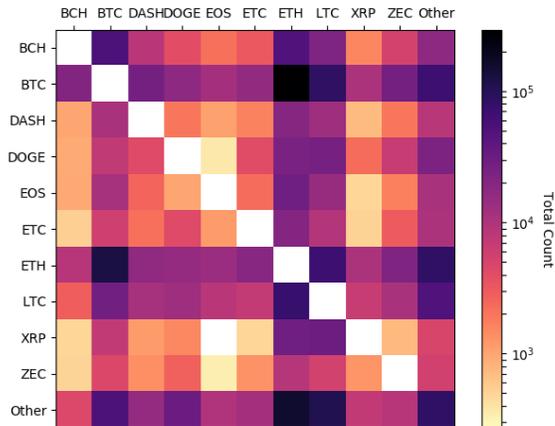


Fig. 2: For each currency pair X-Y, the count (in logarithm scale) of the number of scraped ShapeShift transactions with X as curln and Y as curOut.

For each of these currencies, we ran a full node in order to download the entire blockchain. For the ones supported by the BlockSci tool [4], we used it to parse and analyze their blockchains. BlockSci does not, however, support Ethereum (Classic), Dogecoin, or Bitcoin Cash. For these we thus parsed the blockchains using Python scripts, stored the data as Apache Spark parquet files, and analyzed them using custom scripts. In total, we ended up working with 582.5 GB of raw blockchain data and 255.7 GB of parsed blockchain data.

V. CLUSTERING ANALYSIS

Beyond the information advertised by trading platforms (which in Changelly’s case is effectively nothing), it is also possible to gain insight into their behavior using the same methods that have been developed for Bitcoin-based services [9]; i.e., by interacting with them ourselves and observing which addresses they use. In cryptocurrencies where address clustering is possible, these hand-collected tags can then be combined with cluster data to identify larger volumes of activity associated with the service.

As described in Sections IV-A and IV-B, we engaged in transactions with both ShapeShift and Changelly, focusing on the eight cryptocurrencies for which we had blockchain data (identified in Section IV-C). The one exception was for Dash, which was not supported by Changelly (but which we still interacted with using ShapeShift).

For these blockchains, we then ran the standard “multi-input” clustering heuristic [16], [9], which states that in a transaction with multiple input addresses, all inputs belong to the same entity. Intuitively, this is because the entity in control of a transaction must know the private key associated with any address used as input. While there are techniques today that can be used to invalidate this heuristic, such as CoinJoin [7], it is generally still believed to be safe, especially when applied to services such as exchanges.

We clustered addresses according to this heuristic for five of our blockchains. First, we excluded Bitcoin Cash, which overwhelmed our computational resources every time we tried

to perform the clustering. We also excluded Ethereum and Ethereum Classic, which both use an account-based rather than a UTXO-based model (meaning clustering has no effect). For blockchains with additional functionality, such as Zcash, we ignored their opt-in privacy features and focused on only the Bitcoin-like components.

A. Cluster statistics

Statistics about these clusters can be found in Table II. Looking at this table, it is clear that our results were fairly uneven: for Dogecoin, for example, the three ShapeShift transactions we performed resulted in finding only three addresses, which each had done a very small number of transactions. The three Changelly transactions we performed, in contrast, resulted in us finding 24,893 addresses, which in total had received over 67 trillion DOGE. Similarly, while the clustering was quite effective for both Bitcoin and Zcash it was again rather lopsided for Litecoin, and for both Ethereum and Ethereum Classic we found quite different levels of activity for the addresses we used (transacting in millions of ETH for ShapeShift but far less for Changelly, and millions of ETC for Changelly but orders of magnitude fewer for ShapeShift).

The inconsistency of these results suggests that both of the trading platforms operate a number of different clusters in each cryptocurrency, and perhaps even change their behavior depending on the currency. This makes it difficult to say based on these interactions which of the two services had a higher volume of transactions, as it is clear that we did not capture a comprehensive view of the activity of either. We thus focus more on what the clustering analysis can tell us in terms of the interactions these services had with other services in the cryptocurrency ecosystem (as we do below), and on how it can be used to augment the remainder of our analysis throughout the rest of the paper. In particular we use it to augment our heuristics for identifying on-chain transactions in Section VI, and use it in several case studies that we explore in Section VIII.

B. Cluster interactions

These trading platforms do not exist in a vacuum, meaning we expect them to engage in interactions with other types of services in each of the cryptocurrencies that they support. The available services depend on the cryptocurrency in question, but may range from exchanges to so-called darknet markets to gambling services.

To identify these interactions, we looked at all clusters that had either sent money to or received money from the ShapeShift and Changelly clusters. We focused our attention here solely on Bitcoin and Zcash, as these were the only cryptocurrencies with meaningful clusters. To identify the services associated with these clusters, we used two sets of tagging data: for Bitcoin we used the data available from WalletExplorer,¹⁰ which covers a wide variety of different Bitcoin-based services, and for Zcash we used hand-collected data from Kappos et al. [5], which covers only exchanges.

1) *Zcash*: In Zcash, there were 20,497 clusters that had interacted with Changelly, and 63,808 that had interacted with ShapeShift. Of these clusters, 2610 had interacted with both services. If we ranked the clusters by the total number of coins transacted (in terms of both coins sent and coins received), then for Changelly the top ten clusters accounted for 54% of all the value transacted; for ShapeShift it was only 25%. We present the interactions of these top ten clusters in Table III.

These results demonstrate that many of the tagged clusters belonged to exchanges, but that the majority of them were in fact not tagged. For Changelly, for example, cluster #34097 was the one that had sent and received the second highest number of coins (the equivalent of roughly 48 million USD at the time they were transacted), despite consisting of only four addresses. Besides this anomaly, however, many of the biggest clusters interacting with these trading platforms belonged to exchanges. We observed interactions with exchanges outside of the top ten clusters as well; e.g., Kraken had sent 231 ZEC to ShapeShift and 617 ZEC to Changelly, and Binance had sent 153 ZEC to Changelly. Finally, the dominance of the Poloniex exchange is notable in that it is itself only used for exchanging cryptocurrencies (i.e., it does not support fiat currencies), so these trading platforms do not provide any separate functionality. It is thus unclear why its users would interact with them, other than to attempt to obscure the flow of their coins or to exploit arbitrage opportunities.

2) *Bitcoin*: In Bitcoin, there were 114,151 clusters that had interacted with Changelly, and 37,433 that had interacted with ShapeShift. Of these clusters, 1624 had interacted with both. If we ranked the clusters by the total number of coins transacted (in terms of both coins sent and coins received), then for Changelly the top ten clusters accounted for 66% of all the value transacted and for ShapeShift it was 85%. We present the interactions of these top ten clusters in Table IV.

As for Zcash, these results demonstrate that many of the tagged clusters belonged to exchanges but the majority of them were not tagged. For Changelly, for example, the top cluster consisted of three addresses and had sent over 212K BTC, with one of the addresses having received over 1 billion USD worth of Bitcoin. Based on a manual examination, however, we believe that these top clusters may in fact belong to the respective services themselves (i.e., to ShapeShift and Changelly), and just represent addresses that were not captured by the multi-input heuristic.

Otherwise, these results re-assert the dominance of Poloniex, Bittrex, and HitBTC, and in particular the extent to which their popularity extends across multiple cryptocurrencies rather than being restricted to a single one. Other exchanges which were not included in the top ten clusters were Bitstamp (78 BTC sent to Changelly, 7,052 BTC received from it, and 39 BTC sent to ShapeShift), btc.de (680 BTC sent to ShapeShift), and Huobi (4,315 BTC received from Changelly).

Given that Bitcoin has a more varied ecosystem of different services, it is possible to look beyond exchanges and general information about clusters, and look in addition at the interaction with different types of services. In order to identify all the Bitcoin-based services interacting with ShapeShift and Changelly, we first sorted the tags we had into categories. Some of these were already sorted, and the rest we did

¹⁰<https://www.walletexplorer.com/>

Currency	ShapeShift				Changelly			
	# our txs	# addresses	# txs	# coins received	# our txs	# addresses	# txs	# coins received
BTC	4	72,002	103,292	169,079	4	189,926	213,089	403,893
DASH	4	8	4	0.04	-	-	-	-
DOGE	3	3	8	629	3	24,893	380,457	67,266,293,558
ETH	5	5	2,320,878	8,996,178	9	9	174,836	227,523
ETC	3	2	32	68	5	3	23,730	4,638,390
LTC	4	2009	3527	35,542	4	4	141,981	5,024,460
ZEC	5	81,675	146,417	1,936,966	8	27,685	86,002	380,565

TABLE II: The ShapeShift and Changelly clusters discovered from our personal interactions. The first column in each counts all transactions we did using that currency as either curln or curOut, and the remaining three count the number of addresses, transactions, and received coins associated with the cluster over its entire existence.

ShapeShift		Changelly	
Cluster	# coins	Cluster	# coins
Poloniex	168,615	Poloniex	69,963
Bitfinex	165,979	<i>#34097</i>	68,273
<i>#23</i>	17,535	HitBTC	48,101
Binance	12,791	Bittrex	21,262
<i>#44</i>	10,585	<i>#3773</i>	16,197
<i>#195</i>	2915	<i>#45437</i>	13,903
<i>#4</i>	2561	<i>#21</i>	12,235
<i>#838</i>	1738	<i>#4</i>	7571
<i>#11204</i>	1717	<i>#39</i>	6629
<i>#288734</i>	1451	<i>#413911</i>	6292

TABLE III: The top ten clusters, in terms of total coins transacted, that interacted with the ShapeShift and Changelly clusters in Zcash. Untagged clusters are in italics, and a smaller cluster identifier indicates a larger cluster (in terms of the number of addresses).

ShapeShift		Changelly	
Cluster	# coins	Cluster	# coins
<i>#U1S</i>	78,458	<i>#U1C</i>	223,053
Poloniex	44,038	HitBTC	139,959
Bittrex	31,096	<i>#U1S</i>	38,296
<i>#U4S</i>	15,444	Bittrex	22,701
<i>#U5S</i>	2932	Poloniex	19,743
<i>#U6S</i>	1736	<i>#U6C</i>	15,645
<i>#U7S</i>	1699	<i>#U7C</i>	11,484
<i>#U8S</i>	1666	<i>#U8C</i>	8877
<i>#U9S</i>	1666	<i>#U9C</i>	8227
<i>#U10S</i>	1262	<i>#U10C</i>	7892

TABLE IV: The top ten clusters, in terms of total coins transacted, that interacted with the ShapeShift and Changelly clusters in Bitcoin. Untagged clusters are in italics, using short unique identifiers.

manually. We labelled an address with some associated tag as an *exchange* if it allowed a user to trade coins for some other asset; as a *faucet* if it gave out free coins; as *gambling* if it allowed users to bet their coins for a chance at some reward; as a *merchant* if it offered some good or service; as a *person* if it was specific to an individual; as a *pool* if it belonged to a mining pool; as *unknown* if it did not have an associated tag; and as a *wallet* if it belonged to a wallet provider. Table V shows the results of splitting the clusters interacting with ShapeShift and Changelly into these categories.

In total, the category that had interacted the most with both clusters was (perhaps unsurprisingly) exchanges, although collectively there were in general more coins transacted with clusters belonging to unknown entities. For some of the services, the asymmetry of the interactions is exactly what we

would expect; e.g., you would expect to see money flowing out of faucets and mining pools, but not going back, which explains why there are no transactions into mining pool clusters and very few into faucet ones. Interestingly, the category with one of the lowest volumes of interaction (in terms of either number of transactions or number of coins transacted) was gambling services, which for ShapeShift were even less common than faucets. This, coupled with the fact that we observed no interactions with darknet markets, suggests that either criminals are not using these trading platforms, or that if they are they are being careful in how they do so (by, for example, first moving their illicit bitcoins into a personal wallet before sending them to ShapeShift).

VI. IDENTIFYING BLOCKCHAIN TRANSACTIONS

While there are already interesting insights to be gained from looking just at the data offered by trading platforms or their broader interactions with other cryptocurrency services, in order to gain deeper insights about the way they are used it is necessary to identify not just their internal transactions but also the transactions that appear on the blockchains of the traded currencies. This section presents heuristics for identifying these on-chain transactions, and the next section explores the additional insights these transactions can offer.

Given the clusters identified in Section V, it is already possible to identify on-chain transactions, in terms of the transactions that send money to and take money from these clusters. This is useful for macro-level statistics about the usage of these trading platforms, but doesn't allow us to identify the coins associated with individual ShapeShift transactions. To do this, we must instead identify the two on-chain transactions that are triggered by every transaction made by a user on one of these platforms. In particular, as described in Section III-B, an interaction results in the deposit of coins on the curln blockchain (which we refer to as "Phase 1"), and the withdrawal of coins on the curOut blockchain ("Phase 2").

For both phases, we describe below heuristics for identifying these two types of transactions. Across both, we consider three main requirements: (1) that the candidate transaction occurred reasonably close (in time) to the point at which it was advertised; (2) that the value it carries is reasonably close to the expected amount; and (3) to avoid false positives, that there are no other candidate transactions satisfying these first requirements. We discuss concrete choices for these "reasonable" parameters below, as well as other tweaks necessary

Category	ShapeShift (Bitcoin)				Changelly (Bitcoin)			
	Coins in	# txs	Coins out	# txs	Coins in	# txs	Coins out	# txs
Exchange	62,069	135,049	14,092	11,944	139,885	257,337	57,730	68,236
Faucet	40	83	1	2	1	1	30	108
Gambling	0.23	32	2	38	0.08	21	62	884
Merchant	0.74	39	15	79	0.20	59	102	472
Person	25	821	5	163	54	7665	4041	3172
Pool	2	58	0	0	0.01	1	0	0
Unknown	90,334	222,964	52,934	140,080	306,765	197,014	339,286	694,890
Wallet	35	261	16	179	65	1038	1324	6915

TABLE V: The categories of clusters that sent coins and received coins from the ShapeShift and Changelly clusters in Bitcoin, according to how much they sent and received and how many transactions of each type they engaged in.

for this basic heuristic. The results for both phases, using the optimal parameters, can be found in Table VI.

A. Phase 1

In Phase 1, we seek to identify the deposit transaction on the input (curln) blockchain. For the two requirements (timing and amount) outlined above, we consider the following concrete choices:

Timing: Given transaction timestamp t , we first find the block b (at some height h) on the curln blockchain that was mined at the time closest to t . We then look at the transactions in not only b but at heights $[h - \delta_b, h + \delta_a]$, where δ_b and δ_a are parameters specific to curln. We chose to look at both earlier and later blocks based on the observation in our own transactions of the timestamp published by ShapeShift (which would sometimes be earlier and sometimes later than the on-chain transaction).

Amount: Since the rate is not taken into account for the deposit transaction (and neither is the miner fee), we consider only transactions in which the amount sent to one of the output addresses exactly matches the advertised amount.

For each of our eight currencies, we ran this heuristic for every ShapeShift transaction using curln as the currency in question, with every possible combination of δ_b and δ_a ranging from 0 to 30. This resulted in a set of candidate transactions with zero hits (meaning no matching transactions were found), a single hit, or multiple hits. To rule out false positives, we initially considered as successful only ShapeShift transactions with a single candidate on-chain transaction, although we describe below another heuristic for doing this. We then used the values of δ_b and δ_a that maximized the number of single-hit transactions. As seen in Table VI, the optimal choice of these parameters varies significantly across currencies, according to their different block rates (typically we needed to look further before or after for currencies in which blocks were produced more frequently).

B. Phase 2

In Phase 2, we seek to identify the transaction on the curOut blockchain, in which money is sent to the user. Again, we consider the following concrete choices:

Timing: As with Phase 1, we first find the closest block b to the timestamp t , at some height h , and look for transactions in the blocks at height $[h - \delta_b, h + \delta_a]$.

Currency	Phase 1			Phase 2		
	δ_b	δ_a	% matching	δ_a	error (%)	% matching
BTC	0	1	65.76	0	0.15	25.85
BCH	9	4	73.83	9	0.1	27.92
DASH	5	5	85.09	11	0.4	36.77
DOGE	1	4	74.95	6	0.6	54.05
ETH	5	0	70.66	23	0.4	28.54
ETC	5	0	72.65	15	1	43.72
LTC	1	2	71.88	3	0.3	33.03
ZEC	1	3	87.13	5	0.6	45.96

TABLE VI: For the selected (optimal) parameters, the percentage of ShapeShift transactions for which we found matching on-chain transactions, for both Phase 1 and 2. For Phase 2, we always use $\delta_b = 0$.

Amount: In theory, for given advertised values amt, rate, and fee, the amount sent should be $\text{amt} \cdot \text{rate} - \text{fee}$. To allow for some variation (which we observed in our own transactions), we consider as candidates all transactions carrying a total value within a reasonable error rate of this amount.

Based on our experience transacting with ShapeShift, we always used $\delta_b = 0$, as we observed that they always publish a transaction before paying the client; this makes sense, as they are likely to wait for the deposit transaction to receive some confirmations in the curln blockchain. For δ_a and the error rate, we again “brute-forced” the choice of these parameters by trying all possible options, ranging from 0 to 30 for δ_a and from 0% to 1% for the error rate (in increments of 0.1%). The error rate for each currency is related to the volatility of the exchange rates between this currency and others.

C. An augmented heuristic

While the basic results demonstrate that these heuristics are already quite effective, there is nevertheless a lot of room to improve: especially for Phase 1, in which the amount should match exactly, it should be possible to identify the on-chain transactions in every case. The exception, however, is when the amount transacted is a common unit of currency, such as 0.1. In these cases, there may be many transactions in the interval we examine carrying that value, which means we cannot necessarily isolate which one is the ShapeShift deposit. This explains why our success rate is lowest in Bitcoin, which has the highest volume of transactions (even though we look in only a single block).

To improve our heuristics, we therefore consider other ways to isolate ShapeShift transactions, so that even if there are multiple hits we can still avoid false positives. In particular, we considered an *augmented* heuristic in which we used the clusters we had already associated with ShapeShift in Section V. In Phase 1, if we had multiple hits but only a single one of these transactions had an output address belonging to the ShapeShift cluster, then we considered this the correct transaction. Phase 2 was the same, but we required instead that the input user was ShapeShift.

We ran this augmented heuristic for the five currencies for which we had cluster data. For Dogecoin and Dash, it made a negligible difference, and for Litecoin and Bitcoin it made a small difference as well: for Litecoin the percentage captured in Phase 1 increased to 72% and to 33.1% in Phase 2, and for Bitcoin it increased to 66.1% in Phase 1 and 25.93% in Phase 2. We thus observed a meaningful difference only in Zcash, in which the percentage captured increased to 90.36% in Phase 1 and to 52.18% in Phase 2. This makes sense given that Zcash was the currency for which clustering was the most effective. Given the modest difference overall, however, we continue in future sections to work with only the “normal” version of our heuristic.

VII. TRACKING SHAPESHIFT ACTIVITY

In the previous section, we saw that it was possible in many cases to identify the on-chain transactions associated with the transactions advertised by ShapeShift. In this section, we take this a step further and show how linking these transactions can be used to identify more complex patterns of behavior.

As shown in Figure 3, we consider these for three main types of transactions. In particular, we look at (1) *pass-through* transactions, which represent the full flow of money as it moves from one currency to the other via the deposit and withdrawal transactions; (2) *U-turns*, in which a user who has shifted into one currency immediately shifts back; and (3) *round-trip* transactions, which are essentially a combination of the first two and follows a user’s flow of money as it moves from one currency to another and then back. Our interest in these particular patterns of behavior is largely based on the role they play in tracking money as it moves across the ledgers of different cryptocurrencies. In particular, our goal is to test the validity of the implicit assumption that the use of ShapeShift provides additional anonymity guarantees beyond simply transacting in a given currency.

In more detail, identifying pass-through transactions allows us to create a link between the input address(es) in the deposit transaction in the curIn blockchain and the output address(es) in the withdrawal transaction in the curOut blockchain.

Identifying U-turns allows us to see when a user has interacted with ShapeShift not because they are interested in holding units of the curOut cryptocurrency, but because they see other benefits in shifting coins back and forth. There are several possible motivations for this: for example, traders may shift quickly back and forth between two different cryptocurrencies in order to profit from differences in their price. We investigate this possibility in Section VIII-C. Equally, people performing money laundering or otherwise holding “dirty” money may engage in such behavior under the belief that once

the coins are moved back into the curIn blockchain, they are “clean” after moving through ShapeShift regardless of what happened with the coins in the curOut blockchain.

Finally, identifying round-trip transactions allows us to create a link between the input address(es) in the deposit transaction in the curIn blockchain with the output address(es) in the later withdrawal transaction in the curIn blockchain, which arguably has a more significant impact on the anonymity of a user than the link exposed by identifying pass-through transactions. Again, there are many reasons why users might engage in such behavior, including the trading and money laundering examples given above. As another example, if a curIn user wanted to make an anonymous payment to another curIn user, they might attempt to do so via a round-trip transaction (using the address of the other user in the second pass-through transaction), under the same assumption that ShapeShift would sever the link between their two addresses.

A. Pass-through transactions

Given a ShapeShift transaction from curIn to curOut, we ran both Phase 1 and Phase 2 for this transaction, which allowed us to identify the deposit in the curIn blockchain and withdrawal in the curOut blockchain. If we were successful in both phases in identifying the on-chain transaction, this allowed us to identify the entire pass-through transaction, as depicted in Figure 3a. As discussed above, this has the effect on anonymity of tracing the flow of funds across this ShapeShift transaction and linking its two endpoints; i.e., the input address(es) in the curIn blockchain with the output address(es) in the curOut blockchain. The results, in terms of the raw numbers of transactions, are in Table VII, and in terms of percentages of all possible transactions are in Figure 4.

Both the table and the figure demonstrate that our success in identifying these types of transactions varied, but in general was (as perhaps expected) the probability of finding both the Phase 1 transaction for curIn and the Phase 2 transaction for curOut (meaning the percentages are roughly the product of the percentages given in Table VI). In general, we were least successful at identifying ShapeShift transaction in blockchains with higher volumes of transactions, such as Bitcoin and Ethereum. This was because the number of potential candidates increased significantly as we relaxed the error rate and block range, so we frequently ended up with multiple hits. Even in the worst scenario of Phase 2 in Bitcoin, however, we were still able to identify more than 25% of the total ShapeShift transactions. In total, across all eight currencies we were able to follow the path of 281,113 ShapeShift transactions.

B. U-turns

As depicted in Figure 3b, we consider a U-turn to be a pattern in which a user has just sent money from curIn to curOut, only to turn around and go immediately back to curIn. This means linking two transactions: first, the Phase 2 transaction used to send money to curOut, and then the Phase 1 transaction used to send money back to curIn. In terms of timing and amount, we require that the second transaction happens within 30 minutes of the first, and that it carries within 0.5% of the value of the first transaction (minus the miner fee).

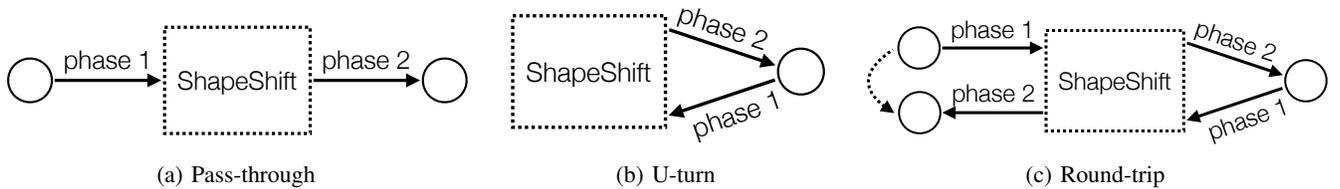


Fig. 3: The different types of transactions we identify, according to how they interact with ShapeShift and which phases are required to identify them.

Currencies	BTC	BCH	DASH	DOGE	ETH	ETC	LTC	ZEC
BTC	-	9596 / 53410	2406 / 10834	2350 / 7427	22237 / 124665	1000 / 5830	5723 / 28711	1350 / 4720
BCH	3955 / 21264	-	249 / 988	268 / 930	1705 / 8867	79 / 545	622 / 2860	135 / 502
DASH	5986 / 27101	1706 / 8358	-	2289 / 4271	4545 / 16804	759 / 2175	3272 / 11524	560 / 1272
DOGE	3080 / 17121	856 / 3984	675 / 1930	-	4118 / 15740	2041 / 4283	4327 / 13138	1114 / 2790
ETH	55173 / 293997	10279 / 50208	4674 / 19881	9824 / 23951	-	5893 / 20648	18237 / 76190	2614 / 8984
ETC	2950 / 16158	629 / 3176	435 / 1676	1246 / 4259	3394 / 19447	-	1438 / 69999	441 / 1322
LTC	17270 / 87613	4892 / 22897	3291 / 13106	9878 / 25871	15102 / 70867	2240 / 9278	-	1843 / 5451
ZEC	6060 / 26790	1329 / 5354	850 / 1943	3123 / 6717	5970 / 21991	1301 / 2998	3734 / 10878	-

TABLE VII: For each pair of currencies, the number of transactions we identified as being a pass-through from one to the other, out of the total number of transactions between those two currencies.

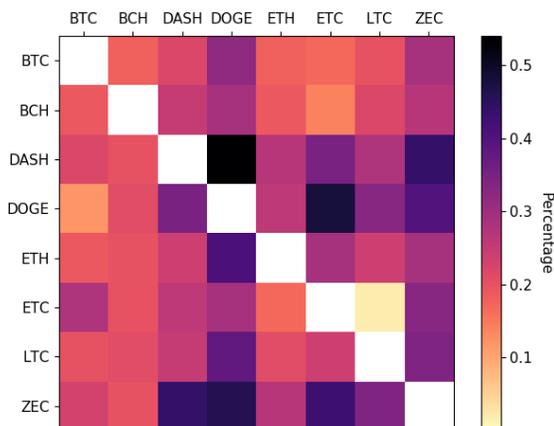


Fig. 4: For each pair of currencies, the number of transactions we identified as being a pass-through from one to the other, as a percentage of the total number of transactions between those two currencies.

While the close timing and amount already give some indication that these two transactions are linked, it is of course possible that this is a coincidence and they were in fact carried out by different users. In order to gain additional confidence that it was the same user, we have two options. In cryptocurrencies based on UTXOs (or *unspent transaction outputs*), each output in a transaction is associated with a new uniquely identifiable UTXO, meaning one address could be associated with potentially many UTXOs. To see if a user is spending a coin immediately, we could thus see if they are spending it from the exact same UTXO that was created in the Phase 2 transaction. In cryptocurrencies based instead on accounts, such as Ethereum, we have no choice but to look just at the addresses. Here we thus define a U-turn as seeing if the address that was used as the output in the Phase 2 transaction

Currency	# (basic)	# (addr)	# (utxo)
BTC	14,669	483	311
BCH	949	55	52
DASH	1590	1182	205
DOGE	390	106	88
ETH	23,407	1560	-
ETC	498	83	-
LTC	4354	894	309
ZEC	496	340	288

TABLE VIII: The number of U-turns identified for each cryptocurrency, according to our basic heuristic concerning timing and value, and both the address-based and UTXO-based heuristics concerning identical ownership. Since Ethereum and Ethereum Classic are account-based, the UTXO heuristic cannot be applied to them.

is also used as the input in the later Phase 1 transaction.

Once we identified such candidate pairs of transactions (tx_1, tx_2) , we then ran Phases 1 and 2, as described in Section VI. In particular, we ran Phase 2 for tx_1 to identify the relevant address in the curOut blockchain and Phase 1 for tx_2 to, again, identify the relevant address in the curOut blockchain.

In fact though, what we really identified in Phase 2 was not an address but, as described above, a newly created UTXO. If the input used in tx_2 was this same UTXO, then we found a U-turn according to the first heuristic. If instead it corresponded just to the same address, then we found a U-turn according to the second heuristic. The results of both of these heuristics, in addition to the basic identification of U-turns according to the timing and amount, can be found in Table VIII.

In total, we identified 46,353 U-turns according to our basic heuristic, 4703 U-turn transactions according to our address-based heuristic, and 2896 U-turn transactions according to our (stricter) UTXO-based heuristic. Additionally, we observe that both Dash and Zcash have been used extensively as “mixer coins” in U-turns, as they hold, respectively, the fourth and

the seventh rank. This rank is higher than their overall rank in terms of popularity, which suggests that users may prefer to use privacy coins as the mixing intermediary under the assumption that this provides some extra degree of anonymity. As the results show, however, Zcash has the highest percentage of identified U-turn transactions. Thus, these users not only do not gain any extra anonymity, but in fact are easily identifiable given that they did not change the address used in 340 out of 496 (68.5%) cases, or—even worse—immediately shifted back the exact same coin they received in 288 (58%) cases.

In the case of Dash, the results suggest something a bit different. In particular, once more, the usage of a privacy coin as the mixing intermediate was not very successful since in 1182 out of the 1590 cases the address which received the fresh coins was the same as the one which shifted it back. It was the exact same coin in only 205 cases, however, which suggests that although the user is the same, there is a local Dash transaction between the two ShapeShift transactions. We defer a further discussion of this asymmetry to Section VIII-B, where we also discuss more generally the use of anonymity features in both Zcash and Dash.

C. Round-trip transactions

As depicted in Figure 3c, a round-trip transaction requires performing two ShapeShift transactions: one out of the initial currency and one back into it. If these transactions happen within a short period of time from each other, we would expect the amount received in the curln blockchain to be fairly close to the initial amount sent. The amounts are not identical, however, because (1) the user has to pay two different miner fees, one for each transaction, and (2) the rate between curln and curOut is likely to change somewhat between the two transactions.

To identify round-trip transactions, we effectively combine the results of the pass-through and U-turn transactions; i.e., we tagged something as a round-trip transaction if the output of a pass-through transaction from X to Y was identified as being involved in a U-turn transaction, which was itself linked to a later pass-through transaction from Y to X (of roughly the same amount). As described at the beginning of the section, this has the powerful effect of creating a link between the sender and recipient within a single currency, despite the fact that money flowed into a different currency in between.

In more detail, we looked for consecutive ShapeShift transactions where for a given pair of cryptocurrencies X and Y: (1) the first transaction was of the form X-Y; (2) the second transaction was of the form Y-X; (3) the second transaction happened relatively soon after the first one; and (4) the value carried by the two transaction was approximately the same. For the third property, we required that the second transaction happened within 30 minutes of the first. For the fourth property, we required that if the first transaction carried x units of curln then the second transaction carried within 0.5% of x -rate units of curOut (using rate from the first transaction). Implicitly, this requires that the interaction in curOut was a U-turn transaction.

As with U-turns, we considered an additional restriction to capture the case in which the user in the curln blockchain stayed the same, meaning money clearly did not change hands. Unlike with U-turns, however, this restriction is less to provide

Currency	# (regular)	# (same addr)
BTC	27,163	58
BCH	665	3
DASH	1262	718
DOGE	235	2
ETH	15,973	9288
ETC	267	90
LTC	4805	1602
ZEC	121	12

TABLE IX: The number of regular round-trip transactions identified for each cryptocurrency, and the number that use the same initial and final address.

safety for the basic heuristic and more to isolate the behavior of people engaged in day trading or money laundering. To identify this pattern, we identify the input addresses used in Phase 1 for the first transaction, which represents the user who initiated the round-trip transaction in the curln blockchain. We then identify the output addresses used in Phase 2 for the second transaction, which represents the user who was the final recipient of the funds. If the address was the same, then it is clear that money has not changed hands. Otherwise, the round-trip transaction acts as a heuristic for linking together the input and output addresses.

The results of running this heuristic (with and without the extra restriction) are in Table IX. In total, we identified 50,491 according to our regular heuristic, and identified 11,773 out of these where the input and output addresses were the same. Across different currencies, however, there was a high level of variance in the results. While this could be a result of the different levels of accuracy in Phase 1 and Phase 2 for different currencies, the heuristic was the same across all of them so the more likely explanation is that users indeed engage in different patterns of behavior with different currencies. For Bitcoin, for example, there was a very small percentage (0.2%) of round-trip transactions that used the same address. This suggests that either users are aware of the general lack of anonymity in the basic Bitcoin protocol and use ShapeShift to make anonymous payments, or that if they do use round-trip transactions as a form of money laundering they are at least careful enough to change their addresses.

In other currencies, however, such as Dash, Ethereum, and Litecoin, there were high percentages of round-trip transactions that used the same input and output address: 57%, 58%, and 33% respectively. In Ethereum, this may be explained by the account-based nature of the currency, which means that it is common for one entity to use only one address. In Dash, as we have already seen in Section VII-B and explore further in Section VIII-B, it may simply be the case that users assume they achieve anonymity just through the use of a privacy coin, so do not take extra measures to hide their identity.

VIII. CASE STUDIES

In this section, we examine potential applications of our analysis, in terms of identifying specific activities in and usages of ShapeShift. As before, our focus is on anonymity, and the potential that such platforms may offer for money laundering or other illicit purposes. To this end, we begin by looking at a case study of coins that were reportedly stolen and mixed by Starscape Capital [17]. We then explore interactions

with so-called privacy coins, in order to understand whether or not ShapeShift users exploit their opt-in anonymity features, followed by ways to identify automated interactions with ShapeShift that may be indicative of day trading, in order to separate this type of behavior from other uses of the platform. Finally, we look at the interactions between ShapeShift and more traditional exchanges, to understand the extent to which unregistered ShapeShift users may nevertheless have their real-world identity revealed.

A. Tracing Starscape Capital’s stolen coins

A new investment firm called Starscape Capital promised investors a 50% return if they invested in their cryptocurrency arbitrage fund. In January 2018 they raised over 2000 ETH (worth 2.2M USD at the time) during their Initial Coin Offering. Shortly afterwards, all of their social media accounts disappeared, and it was reported that an amount of ETH worth \$517,000 was sent from their wallet to ShapeShift, where it was shifted into Monero [17].

Indeed, we observed that the address owned by Starscape Capital participated in 192 transactions across a three-day span, during which it received and sent 2,038 ETH. Of the 133 transactions sent, 109 were to ShapeShift, and 103 were shifts to Monero (the remaining 6 were shifts to Ethereum). The total amount shifted into Monero was 465.61803925 ETH.

Furthermore, within a three-day span after the initial shifts we identified 6 ShapeShift transactions whose carried value was roughly equivalent to the value shifted into Monero, and which exchanged Monero to other coins. This indicates that the coins may have been sent back through ShapeShift, although here of course a more detailed analysis would be required.

B. Usage of anonymity tools

Given the potential usage of ShapeShift for money laundering or other criminal activities, we sought to understand the extent to which its users seemed highly motivated to hide the source of their funds. While using ShapeShift is already one attempt at doing this, we focus here on the combination of using ShapeShift and so-called “privacy coins” that are designed to offer improved anonymity guarantees. Of the blockchain data we had, this meant looking at Dash and Zcash.

Before diving into the anonymity features of these coins, it is worth noting that—as we saw in Table VIII—a non-trivial fraction of the transactions into these currencies moved their money right back out: 5% in Dash and 58% in Zcash according to our stricter heuristic, and 68.3% in Dash and 69% in Zcash according to our looser one. While this is perhaps already an indication that many of the users interacting with these currencies were not interested in their anonymity features, we nevertheless examined whether or not the remaining transactions did exploit them.

1) *Zcash*: The main anonymity feature in Zcash is known as the *shielded pool*. Briefly, transparent Zcash transactions behave just like Bitcoin transactions in that they reveal in the clear the sender and recipient (according to so-called *t-addresses*), as well as the value being sent. This information is hidden to various degrees, however, when interacting with the pool. In particular, when putting money into the pool the

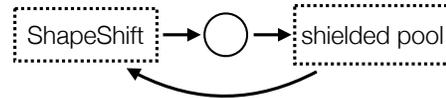


Fig. 5: The two types of interactions we investigated between ShapeShift and the shielded pool in Zcash.

recipient is specified using a so-called *z-address*, which hides the recipient but still reveals the sender, and taking money out of the pool hides the sender (through the use of zero-knowledge proofs [2]) but reveals the recipient.

ShapeShift does not support the use of z-addresses, which means it is not possible for users to shift money directly into the shielded pool. This left two possible interactions for us to investigate, as depicted in Figure 5: (1) a user shifts to a t-address but then uses that t-address to put money into the pool, and (2) a user takes money out of the pool directly into ShapeShift.

For the first type of interaction, we found 14,852 transactions that involved ShapeShift sending ZEC to a t-address (identified in Phase 2). Of these, there were 479 where the next transaction (i.e., the transaction in which this t-address spent its contents) involved putting money into the pool. The total value put into the pool in these transactions was 2089 ZEC, which represented 7% of all the value sent in Phase 2 transactions in Zcash. For the second type of interaction, we found 76,494 transactions that had sent money to ShapeShift (identified in Phase 1). Of these, 2,881 came directly from the pool, with a total value of 12,404 ZEC (representing 16% of the value in Phase 1 transactions).

Thus, while the usage of the anonymity features in Zcash was not necessarily a significant fraction of the overall usage of Zcash in ShapeShift, there is clear potential to move large amounts of Zcash (representing over 10 million USD at the time it was transacted) by combining ShapeShift with the shielded pool. This would make it much harder to follow the flow of money after the Zcash was shifted back into the original currency, even given recent research de-anonymizing certain types of interactions with Zcash’s shielded pool [14], [5] and our results identifying round-trip transactions.

2) *Dash*: As in Zcash, the “standard” transaction in Dash is similar to a Bitcoin transaction in terms of the information it reveals. Its main anonymity feature—*PrivateSend* transactions—are a type of CoinJoin [7]. A CoinJoin is specifically designed to invalidate the multi-input clustering heuristic described in Section V, as it allows multiple users to come together and send coins to different sets of recipients in a single transaction. If each sender sends the same number of coins to their recipient, then it is difficult to determine which input address corresponds to which output address, thus severing the link between an individual sender and recipient.

In a traditional CoinJoin users must find each other in some offline manner (e.g., an IRC channel) and form the transaction together over several rounds of communication.

This can be a cumbersome process, so Dash aims to simplify it for users by automatically finding other users for them and chaining multiple mixes together. In order to

ensure that users cannot accidentally de-anonymize themselves by sending uniquely identifiable values, these PrivateSend transactions are restricted to specific denominations: 0.01, 0.1, 1, and 10 DASH. As observed by Kalodner et al. [4], however, the CoinJoin denominations often contain a fee of 0.0000001 DASH, which must be factored in when searching for these transactions. Our parameters for identifying a CoinJoin were thus that (1) the transaction must have at least three inputs, (2) the outputs must solely consist of values from the list of possible denominations (modulo the fees), and (3) and all output values must be the same.

We first looked to see how often the DASH sent to ShapeShift had originated from a CoinJoin, which meant identifying if the inputs of a Phase 1 transaction were outputs from a CoinJoin. Out of 77,322 single hits we found 1,164 that came from a CoinJoin, carrying a total of 4206 DASH in value (4% of the total value across Dash Phase 1 transactions). Next, we looked at whether or not users performed a CoinJoin after receiving coins from ShapeShift, which meant identifying if the outputs of a Phase 2 transaction had been spent in a CoinJoin. Out of 22,158 single hits we found only 10 CoinJoin transactions, carrying a total of 20 DASH in value (0.07% of the total value across Dash Phase 2 transactions).

If we revisit our results concerning the use of U-turn transactions in Dash from Section VII-B, we recall that there was a large asymmetry in terms of the results of our two heuristics: only 13% of the U-turns used the same UTXO, but 74% of them used the same address. This suggests that some additional on-chain transaction took place between the two ShapeShift transactions, and indeed upon further inspection we identified many cases where this transaction appeared to be a CoinJoin. There thus appears to have been a genuine attempt to take advantage of the privacy that Dash offers, but this was completely ineffective due to the use of the same address to both send and receive the mixed coins.

C. Trading bots using ShapeShift

ShapeShift, like any other cryptocurrency exchange, can be exploited by traders who wish to take advantage of the volatility in cryptocurrency prices. In particular, users can purchase currencies whose price they believe will rise in the future by exchanging them for other currencies whose price they believe will decline. The potential advantages of doing this via ShapeShift, as compared with other platforms that focus more on the exchange between cryptocurrencies and fiat currencies, are that (1) ShapeShift transactions can be easily automated via their API, and (2) a single ShapeShift transaction acts to both purchase desired coins and dump unwanted ones. Such trading usually requires large volumes of transactions and high precision on their the timing, due to the constant fluctuation in cryptocurrency prices.

We thus looked through our scraped transactions for activity that could be identified as being associated with trading bots.

Initially, we searched for sets of consecutive ShapeShift transactions that carried approximately the same value in USD (with an error rate of 1%) and involved the same currencies. When we did this, however, we found thousands of such sets. We thus added the extra conditions that there must be at least

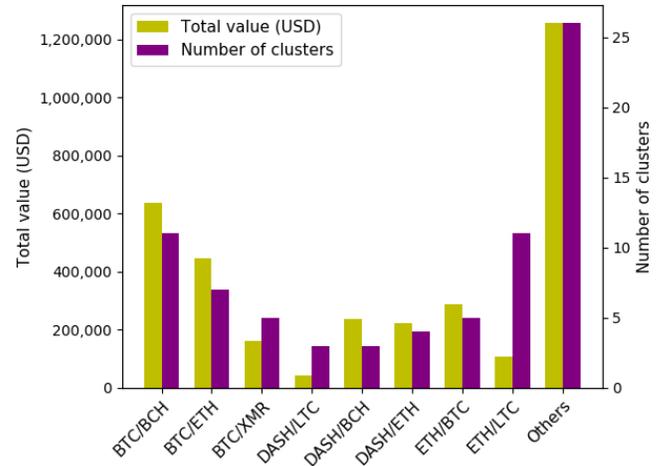


Fig. 6: Our 75 clusters of likely trading bots, categorized by the pair of currencies they trade between and the total amount transacted by those clusters (in USD).

15 transactions in the set that took place in a span of five minutes; i.e., that within a five-minute block of ShapeShift transactions there were at least 15 involving the same currencies and carrying the same approximate USD value. This resulted in 75 such sets. Given that there were roughly 40 ShapeShift transactions within a span of five minutes, these clusters of transactions already represented 37.5% of the total activity within ShapeShift in that period.

After obtaining our 75 clusters, we manually removed transactions that we believed were “false positives” in that they happened to have a similar value but were clearly the odd one out. For example, in a cluster of 20 transactions that involved 19 transactions from ETH to BTC and one from LTC to ZEC, we removed the latter. We were thus left with clusters of either a particular pair (e.g., ETH to BTC) or two pairs where the curOut or the curIn was the same (e.g., ETH to BTC and ZEC to BTC), which suggests either the purchase of a rising coin or the dump of a declining one. We further sought to validate these clusters by using the results of Phases 1 and 2 to determine if the clusters shared common addresses. While our results were less even in UTXO-based currencies (as most entities operate using multiple addresses), in account-based chains there was in almost every case one particular address that was involved in the transactions in the cluster.

We summarize the results of this form of clustering in Figure 6, in terms of the commonality of different pairs of currencies and the total money exchanged by a given cluster. It is clear that the most common interactions are performed between the most popular currencies overall, with the exception of Monero (XMR). In particular, we found five clusters consisting of between 17 and 20 transactions that exchanged BTC for XMR. This may suggest that the entities performing those clusters were interested in the increased anonymity guarantees of Monero, which suggests it may be fruitful to apply the analysis in Section VIII-B to Monero as well.

Beyond these clusters, we also attempted to identify au-

Exchange	# txs	# coins	BTC	LTC	ETH	XMR	DASH
Bitfinex	11	23.4	3.2	0.07	0.09	0	0
Bithumb	3	0.4	0.2	0	0.1	0	0
Bittrex	148	230.7	86.2	23.1	87	5.186	19.3
Changelly	5	11.4	5.7	5.7	0	0	0
Huobi	2	3	0.05	0	2.9	0	0
Kraken	69	60.3	11.2	1.9	20.2	13.8	11.8
Poloniex	83	79.3	40.5	11.5	21.7	2.5	0

TABLE X: The exchanges sending ZEC to ShapeShift, according to the total number of transactions, the total ZEC sent, and the ZEC sent that came from each of five other currencies (Bitcoin, Litecoin, Ethereum, Monero, and Dash).

tomated behavior by looking at the values being transacted. Unsurprisingly, the most common denominations were integers between 0 and 10, or large integers multiple of 100 for currencies with low value. Indeed, among the 500 most common values, 454 had four or fewer decimal points. We found, however, several notable exceptions that suggest automated behavior likely to be associated with trading. In particular, among the 100 most common values we encountered the values 0.013279, 0.31089377, 0.13939789, and 0.997853, each of which appeared in thousands of transactions. Furthermore, each of these four unique values appeared in transactions only where the curln was the same. In particular, all transactions of the first type shift 0.013279 BTC to multiple currencies, of the second type shift 0.13939789 ETH to multiple currencies, and of the third and fourth types shift that amount in LTC to multiple currencies.

D. Exchanges using ShapeShift

As we already observed in Section V, there are many interactions between traditional exchanges and ShapeShift, as users may exchange fiat currencies for “bigger” cryptocurrencies such as Bitcoin and Ethereum and then use ShapeShift to exchange into other cryptocurrencies that cannot be exchanged with fiat currencies. In this section, we revisit this analysis and combine the information about the exchanges that have interacted with ShapeShift with the outputs of Phases 1 and 2 in order to understand the currencies into and out of which these exchange users were shifting. Given that these traditional exchanges comply with KYC/AML regulations and collect information about the real-world identities of their users, users who transact with ShapeShift directly from their exchange accounts may taint these transactions with that information as well, in addition to any transactions they carry out in the shifted currency (assuming we were able to identify the corresponding pass-through transaction in Section VII). Again, we split our analysis according to the two cryptocurrencies for which we had tagging information about exchanges: Zcash and Bitcoin.

1) *Zcash*: As identified by combining the clusters from Section V with the Phase 1 transactions from Section VI, the exchanges sending ZEC to ShapeShift can be seen in Table X. The exchanges receiving ZEC from ShapeShift (according to Phase 2 transactions) can be seen in Table XI. In addition to the total numbers of ZEC transacted with these exchanges, these tables also break them down according to the most popular currencies with which to exchange them.

Looking at both tables, we observe that the number of coins transacted is far lower here than it was in Section V-B1; this is

Exchange	# txs	# coins	BTC	LTC	ETH	XMR	DASH
Binance	156	396.7	115.4	43.5	185	17.8	1.4
Bitfinex	27	62.8	8.5	11.5	17.4	0	0.2
Bithumb	31	60	26	1.6	25	0	0
Bitrix	121	211.9	68.7	22.7	78.9	0.2	12.6
Changelly	19	50.1	1.2	0.5	1.6	1.0	33.4
Exmo	35	13.6	1	2.5	5.7	1.6	0
HitBTC	18	33.5	8.5	0.5	19.6	0	0.2
huobi	13	40.3	7.3	0	2.3	0	3.8
Kraken	46	41	4.2	13.8	2.6	0	1.2
Poloniex	73	90.7	35.5	7.5	21.5	1	1.3

TABLE XI: The exchanges receiving ZEC from ShapeShift, according to the total number of transactions, the total ZEC received, and the ZEC received that came from each of five other currencies (Bitcoin, Litecoin, Ethereum, Monero, and Dash).

Exchange	# txs	# coins	ETH	XMR
Bitstamp	111	11.2	1.6	4
Bittrex	185	17.4	10.8	2.4
Huobi	21	3.1	1.0	1.4
Poloniex	86	4.7	1.7	0

TABLE XII: The exchanges sending BTC to ShapeShift, according to the total number of transactions, the total BTC sent, and the BTC sent that came from each of two other currencies (Ethereum and Monero).

natural, however, as we look here at the transactions conducted over the six-month period for which we have ShapeShift data, whereas the numbers reported in Section V-B1 were over the lifetime over both clusters of addresses (but could not pinpoint information about specific transactions).

For both transactions with Zcash as curln and with Zcash as curOut, we found that the most popular currencies to exchange with were Ethereum, Bitcoin, Litecoin, Dash, and Monero. As in Section VIII-C, these largely follow the currencies that are overall popular in ShapeShift, with the exception of Monero. We also found significantly more transactions with Zcash used as curOut. Interestingly, we observed a number of interactions between Changelly and ShapeShift, even if Changelly was in general used less often than other exchanges. This suggests that users may try to improve their anonymity by interleaving usage of the two trading platforms, or attempt to profit from better exchange rates.

2) *Bitcoin*: As identified by combining the clusters from Section V with the Phase 1 transactions from Section VI, the exchanges sending BTC to ShapeShift can be seen in Table XII. The exchanges receiving BTC from ShapeShift (according to Phase 2 transactions) can be seen in Table XIII. In addition to the total numbers of BTC transacted with these exchanges, these tables also break them down according to the most popular currencies with which to exchange them.

As with Zcash, the number of coins transacted is far lower here than it was in Section V-B2, and we identified far more transactions with Bitcoin used as curOut. There was less variance across the choice of currencies with Bitcoin used as curln, however, with only Ethereum and Monero reflecting any large amount of activity. The only notable exception was for Poloniex, for which the second-most used currency (in terms of the number of coins transactions) was FunFair, to which we saw two transactions totalling 0.8 BTC in value. FunFair is a token associated primarily with gambling that we did not see

Exchange	# txs	# coins	ETH	LTC	XMR	DASH
Bitstamp	543	70.1	10.8	29.7	4.6	5.4
Bittrex	4020	415.6	61.2	193.1	15.4	34.4
HitBtc	195	17	2.9	7.5	0.0	2.1
Huobi	2074	156.8	20.1	78.3	8	7.1
Poloniex	1443	112.7	13.7	55.9	5.6	5.2
Changelly	182	10.9	3.2	4.9	0.3	0.1

TABLE XIII: The exchanges receiving BTC from ShapeShift, according to the total number of transactions, the total BTC received, and the BTC received that came from each of four other currencies (Ethereum, Litecoin, Dash, and Monero).

being used by almost any other exchanges (at least not in any noticeable volume).

IX. CONCLUSIONS

In this study, we presented a characterization of the usage of the ShapeShift trading platform over a six-month period, focusing on the interactions these types of platforms have with other cryptocurrency-based services and on their potential role in enabling money laundering or other forms of criminal activity via their ability to seamlessly link together the ledgers of multiple different cryptocurrencies. To accomplish this task, we looked at these trading platforms from several different perspectives, ranging from the macro-level view of the behavior of their clusters of addresses to the correlations between the transactions they produce in the cryptocurrency ledgers themselves. The techniques we develop demonstrate that it is possible to capture complex transactional behaviors and trace their activity even as it moves across ledgers, which has implications for any criminals attempting to use these platforms to obscure their flow of money.

ACKNOWLEDGMENTS

We would like to thank Bernhard Haslhofer and Rainer Stütz for performing the Bitcoin clustering for us using the GraphSense tool. The authors are supported by the EU H2020 TITANIUM project under grant agreement number 740558.

REFERENCES

- [1] E. Androulaki, G. Karame, M. Roeschlin, T. Scherer, and S. Capkun. Evaluating user privacy in Bitcoin. In A.-R. Sadeghi, editor, *FC 2013*, volume 7859 of *LNCS*, pages 34–51, Okinawa, Japan, Apr. 1–5, 2013. Springer, Heidelberg, Germany.
- [2] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474, Berkeley, CA, USA, May 18–21, 2014. IEEE Computer Society Press.
- [3] J. Dunietz. The Imperfect Crime: How the WannaCry Hackers Could Get Nabbed, Aug. 2017. <https://www.scientificamerican.com/article/the-imperfect-crime-how-the-wannacry-hackers-could-get-nabbed/>.
- [4] H. Kalodner, S. Goldfeder, A. Chator, M. Möser, and A. Narayanan. Blocksci: Design and applications of a blockchain analysis platform, 2017. <https://arxiv.org/pdf/1709.02489.pdf>.
- [5] G. Kappos, H. Yousaf, M. Maller, and S. Meiklejohn. An empirical analysis of anonymity in Zcash. In *Proceedings of the USENIX Security Symposium*, 2018.
- [6] A. Kumar, C. Fischer, S. Tople, and P. Saxena. A traceability analysis of monero’s blockchain. In S. N. Foley, D. Gollmann, and E. Sneekenes, editors, *ESORICS 2017, Part II*, volume 10493 of *LNCS*, pages 153–173, Oslo, Norway, Sept. 11–15, 2017. Springer, Heidelberg, Germany.

- [7] G. Maxwell. Coinjoin: Bitcoin privacy for the real world. In *Post on Bitcoin forum*, 2013.
- [8] S. Meiklejohn and C. Orlandi. Privacy-enhancing overlays in bitcoin. In M. Brenner, N. Christin, B. Johnson, and K. Rohloff, editors, *FC 2015 Workshops*, volume 8976 of *LNCS*, pages 127–141, San Juan, Puerto Rico, Jan. 30, 2015. Springer, Heidelberg, Germany.
- [9] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage. A fistful of bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 Internet Measurement Conference*, pages 127–140. ACM, 2013.
- [10] M. Möser and R. Böhme. Anonymous alone? measuring Bitcoin’s second-generation anonymization techniques. In *IEEE Security & Privacy on the Blockchain (IEEE S&B)*, 2017.
- [11] M. Möser, K. Soska, E. Heilman, K. Lee, H. Heffan, S. Srivastava, K. Hogan, J. Hennessey, A. Miller, A. Narayanan, and N. Christin. An empirical analysis of linkability in the Monero blockchain. *Proceedings on Privacy Enhancing Technologies*, pages 143–163, 2018.
- [12] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. bitcoin.org/bitcoin.pdf.
- [13] R. R. O’Leary. Will Lightning Help or Hurt Bitcoin Privacy?, Feb. 2018. <https://www.coindesk.com/will-lightning-help-hurt-bitcoin-privacy/>.
- [14] J. Quesnelle. On the linkability of Zcash transactions. arXiv:1712.01210, 2017. <https://arxiv.org/pdf/1712.01210.pdf>.
- [15] F. Reid and M. Harrigan. An analysis of anonymity in the Bitcoin system. In *Security and privacy in social networks*, pages 197–223. Springer, 2013.
- [16] D. Ron and A. Shamir. Quantitative analysis of the full Bitcoin transaction graph. In A.-R. Sadeghi, editor, *FC 2013*, volume 7859 of *LNCS*, pages 6–24, Okinawa, Japan, Apr. 1–5, 2013. Springer, Heidelberg, Germany.
- [17] J. Scheck and S. Shifflett. How dirty money disappears into the black hole of cryptocurrency, Sept. 2018. <https://www.wsj.com/articles/how-dirty-money-disappears-into-the-black-hole-of-cryptocurrency-1538149743>.
- [18] M. Spagnuolo, F. Maggi, and S. Zanero. BitIodine: Extracting intelligence from the bitcoin network. In N. Christin and R. Safavi-Naini, editors, *FC 2014*, volume 8437 of *LNCS*, pages 457–468, Christ Church, Barbados, Mar. 3–7, 2014. Springer, Heidelberg, Germany.
- [19] E. Voorhees. Announcing ShapeShift membership, Sept. 2018. <https://info.shapesift.io/blog/2018/09/04/introducing-shapesift-membership/>.