

“Coins for Bombs”

Increased Transparency of the Global Financial System: Evidence from Terrorist Attacks Financing Detection in Blockchain-based Currencies

Dan Amiram* Bjørn N. Jørgensen† Daniel Rabetti‡

April 2021

Abstract

The proliferation of blockchain-based cryptocurrencies, which essentially use public accounting ledgers, has two opposing effects related to the global financial system’s transparency. On the one hand, governments, market regulators, and financial institutions make significant efforts to curtail the financing of illicit activities, and cryptocurrencies may impede these efforts. On the other hand, funds and transfers that used to be known only by the involved parties are now transparent on the public blockchain ledger. This study empirically examines whether outsiders can identify and predict the financing of on-the-ground operations of terrorist attacks on the public blockchain systems. We do so by using empirical strategies based on accounting and finance literature, forensic accounting, and machine learning algorithms. We provide evidence that blockchain-based cryptocurrencies are used to finance on-the-ground operations of terrorist attacks. However, the blockchain ledger underlying transparency also enables outsiders to identify the fund trails and predict attacks.

JEL classification: G15, G18, G29, K29, K42, O16.

Keywords: Transparency, Terrorist Financing, Economics of Blockchain, Accounting, Bitcoin.

We are grateful to Luzi Hail (editor), an anonymous reviewer, Marie Briere (discussant), Brian Burnett, Evgeny Lyandres, Steve Rock, Daniel Scott Cohen, Andrew Wu (discussant), Peter Zimmerman (discussant), and participants of the Crypto and Blockchain Economics Research Conference (CBER), the 3rd Bergen FinTech Conference, the 4th Shanghai-Edinburgh Fintech Conference and the 6th Fintech International Conference for invaluable comments. We thank WhiteStream for providing intel on the Sri Lanka Eastern bombing, and DeepSeek for additional insights about terrorist’s operations in the dark web. Daniel Rabetti is thankful to Tel Aviv University and Hogege Blockchain Research Institute for financial support. All errors are our own.

*danamiram@tauex.tau.ac.il. Coller School of Management, Tel Aviv University.

†bnj.acc@cbs.dk. Copenhagen Business School.

‡rabetti@mail.tau.ac.il. Coller School of Management, Tel Aviv University.

The global financial system, responsible for trillions of dollars of funds transferred worldwide annually, is opaque for outsiders who are not part of the specific transaction. Financial institutions involved in the transfer can observe only certain parameters of it, and regulators can usually observe only information related to their jurisdiction. In recent years, there has been significant growth in international money transfers through blockchain-based cryptocurrencies. Blockchain is an accounting technology as it is concerned with the transfer of ownership of assets, and maintaining a ledger of accurate public financial information (ICAEW (2018)). These ledgers record details of transfers, such as the sender's and receiver's wallet addresses, the amount transferred, and the date and time of the transfer. The proliferation of cryptocurrencies has two opposing effects related to the transparency of the international money-transfer system.¹ On the one hand, governments, financial institutions, and financial market regulators invest hundreds of millions of dollars and numerous person-hours in curtailing illegal international financing through the traditional financial system (Belasco, Eaglen, Hartig, Jonas, McCord, and Mueller (2018)). Cryptocurrencies and in particular, Bitcoin, the most popular blockchain-based cryptocurrency, may enable criminals to circumvent these efforts. On the other hand, fund transfers that used to be known only to the involved parties are now transparent to anyone with the technical knowledge of the blockchain system and with the ability to analyze information in public ledgers that record the transfer of ownership of assets and financial information. The increased transparency of transfers can help outsiders identify and predict illicit activities by monitoring abnormal transaction behavior.

This study empirically examines whether outsiders can identify and predict illegal international activities, specifically the financing of on-the-ground operations of large-scale international terrorist attacks, via the public blockchain system.² By financing on-the-ground attacks, we mean the financing by the terror organization of the actual operation of the attack, such as paying for explosives and weapons, as opposed to funding the terror organization itself by donation campaigns. This distinction is important because the former has short-term characteristics and requires the use of intense money laundering activities, permitting the analysis of Bitcoin abnormal volume as a proxy for terrorist financing in the vicinity of terrorist

¹About 4,800 cryptocurrencies are currently listed in more than 300 crypto exchanges across the globe (see <https://coinmarketcap.com/currencies/bitcoin/>).

²We use the terms "on-the-ground operations of large-scale international terrorist attacks" and just "terrorist attacks" interchangeably in this paper.

attacks. As the data on fund transfers are not publicly available in the traditional global financial system, this is the first time that researchers are able to conduct this type of investigation to the best of our knowledge. We choose terrorist attacks as our setting for illegal financing activities, as terrorist attacks are observable, require meaningful financing and are, jointly with money laundering, activities that the traditional financial system spends significant resources to curtail.

Our study is motivated by several considerations. First and foremost, we, and more generally the literature on transparency and disclosure, are inherently interested in the causes and consequences of the tradeoff between more detailed, privately-disclosed information, and a less detailed but more widely-distributed disclosure which is the tradeoff faced by the terror financier.³ On the one hand, fund transfers are not in the public domain in the traditional banking system. In the traditional banking system, fund transfers are opaque to outsiders including, to a large extent, regulators. Outsiders that are not part of the banking system cannot observe fund transfers at all. Moreover, even regulators generally depend on the information they receive from banks by law or court rulings to gain insights on transfers. Beyond that, the global regulatory system is fragmented, and regulators can generally follow only parts of the fund trail which are in their jurisdiction. Nevertheless, the global banking system generally has significant know-your-customer (KYC) policies in place, and although challenging, it is generally possible for law enforcement agencies to identify the funds' owner. On the other hand, fund transfers on the chain are in the public domain in the Bitcoin blockchain. Further, unless unique privacy techniques are applied, fund transfers can be traced to wallets. Wallet owners are anonymous (although their activity - fund transfers, is public). Except, some wallet owners reveal their identity willingly or by accident. For instance, we focus on users that are business entities in the chain (e.g., crypto exchange). These users willingly provide wallet addresses to their clients; Thus, it is public information that a specific wallet belongs to a specific user in some cases. With the development of algorithms, such as the one we employ in our study, it is possible to map the same user's address by checking whether a known address co-spent with another address on the chain.⁴ At the end of mapping an address, we can account for the proper user balances at a given time and

³See [Beyer, Cohen, Lys, and Walther \(2010\)](#) for a review of some of this literature.

⁴By co-spending, we mean that two addresses or more share the same input in one transaction, which indicates that they must belong to the same user. See more details in section 1.1 below.

the flow of funds throughout time. Moreover, the computer science literature develops refined techniques that allows parties to identify users' IP addresses in some cases (e.g., Kang, Lee, Ko, Woo, and Hong (2020)). Because users' identities may be revealed through these techniques, and transactions are public, Bitcoin is considered pseudo-anonymous.

Therefore, the terror financier's decision heavily depends on a disclosure choice. On the one hand, the financier can choose to reveal information to the traditional system and go through the standard KYC procedures and make all require disclosures privately to the financial institution with which they work. Alternatively, the financier can reveal less detailed information by revealing information on the transactions publicly on the blockchain. The financier essentially trades-off disclosing more detailed information privately against disclosing less information publicly, taking into account the costs of each disclosure path, such as the damage to operational secrecy of the attack before it takes place if this disclosure help outsiders detect the attack ahead of time. In this sense, the terror financier problem is similar to the tradeoff between public and bank debt. The borrower needs to reveal much more detailed information to the bank (e.g., business plans and future projections) than in public debt, which requires less detailed information (e.g., Dhaliwal, Khurana, and Pereira (2011)); however, public debt reveals information to the public, which carries disclosure costs such as revealing proprietary information to competitors and others.

Second, our study is also motivated by the relative advantage that individuals or teams with blockchain knowledge, and (arguably) accounting expertise, have in analyzing blockchain transactions. The Institute of Chartered Accountants in England and Wales, one of the largest and oldest accounting organizations globally describes blockchain as "*an accounting technology. It is concerned with the transfer of ownership of assets, and maintaining a ledger of accurate financial information. The accounting profession is broadly concerned with the measurement and communication of financial information, and the analysis of said information. Much of the profession is concerned with ascertaining or measuring rights and obligations over property, or planning how to best allocate financial resources. For accountants, using blockchain provides clarity over ownership of assets*" (ICAEW (2018)). Under this conceptual view of blockchain, accounting expertise (broadly speaking) is useful since blockchain can be viewed as an accounting system that maintains a ledger of accurate financial information and provides clarity over ownership of assets. As

such, accountants, and accounting researchers may have an advantage as they are experts in understanding the maintenance of a ledger of accurate financial information that provides clarity over ownership of assets. For example, in this study we use accounting and other forensic tools to exploit blockchain transparency and learn about financing activities that we do not have access to in the traditional financial system.

Third, the ability to detect and predict terrorist attacks can hinder terror activities as it puts at risk the operational security that is required for such attacks to succeed. Providing evidence of terror activities using public information, and the transparency of the public blockchain system, increases the probability for greater scrutiny of these activities by the public, researchers, and regulators to prevent such activities. Moreover, finding evidence of terror financing through the public blockchain system suggests that transparency alone, without specific supporting regulation, is not enough to prevent illicit activities. Providing evidence of international terror financing using cryptocurrencies reveals that the unobservable current efforts to counter those activities in the traditional financial system create real barriers for terror financiers; otherwise, they would have little reason to change the way they operate and create new costly setups.

Lastly, beyond the direct physical and psychological damage they cause, terrorist attacks have substantial adverse effects on stock markets and the macro economy. [Eckstein and Tsiddon \(2004\)](#) find that terrorist attacks have short-term impacts on consumption, investment, and net exports. [Abadie and Gardeazabal \(2003\)](#) find that, after the outbreak of terrorism in the late 1960s, per capita GDP in the Basque Country, declined about ten percent relative to a synthetic control region without terrorism. Another consequence is that terrorism is associated with redirecting economic activity away from investment and government spending ([Blomberg, Hess, and Orphanides \(2003\)](#)). The impact of terrorism is also observed in financial markets. Studying the effect of attacks in which publicly-traded firms are targets, [Karolyi and Martell \(2011\)](#) find a stock price reaction of -0.83% at the vicinity of attacks, which corresponds to an average loss per firm per attack of \$401 million in market capitalization. [Arin, Ciferri, and Spagnolo \(2008\)](#) complement this by finding effects on stock market returns and volatility across several equity markets. Given these considerations, our investigation may be useful for accounting and finance researchers, regulators, financial institutions, and other market participants.

We start our investigation by focusing on Bitcoin, as it is the most liquid and used public blockchain-based cryptocurrency.⁵ Its features, such as pseudo-anonymity, wide acceptance, and considerable use in dark markets, make it the most likely cryptocurrency candidate for terrorist organizations to finance their activities. While current evidence indicates that terrorist donation campaigns using Bitcoin raise only marginal amounts, there is no evidence we are aware of whether Bitcoin or any other cryptocurrency is used to finance the operations of on-the-ground terrorist attacks.⁶ Moreover, as opposed to donation campaigns, there is also no evidence of which we are aware that blockchain's inherent transparency enables the ex-post detection of international terrorist activities, and perhaps more importantly, the ex-ante prediction of attacks.

This lack of evidence may result from the transparent nature of the underlying public blockchain ledger system of cryptocurrencies. Bitcoin transparency can potentially mitigate terrorist incentives to use Bitcoin to finance on-the-ground terror attacks, not only because this transparency can help reveal user-identity in some instances, but mainly because it makes it easier for outsiders (including, but not limited to, counter-terror organization) to ex-ante identify and predict future terror attacks. Abnormal Bitcoin volume patterns and anomalous transactions before a large-scale terror attacks, such as we predict in this study, alone, or together with other red flags, such as chatter in radical groups' websites and private "Humint" or "Sigint" intelligence, can threaten the operational secrecy of a terror attack, which is crucial to its success. As we detail above, it is difficult even for law enforcement agencies, especially in an international setting, to get real-time information about on-the-ground-terror financing from current traditional banking. Therefore, as opposed to donation campaigns, the lack of evidence of financing of

⁵Bitcoin recently surpassed 1.2 trillion dollars in total market capitalization, has more than 18 million units in circulation and over half a billion transactions since its inception in 2009 (see Nakamoto (2008) for detailed description of the Bitcoin protocol).

⁶Recently, the U.S. Department of Justice (DOJ) dismantled the Bitcoin campaigns of three terrorist groups: al Qassam Brigades, al Qaeda, and the Islamic State of Iraq and the Levant (ISIS). The action seized cryptocurrencies in over 300 accounts. Chainalysis, a large blockchain analytics provider, reports that two large donation campaigns from designated terrorist groups were able to raise only negligible amounts (<https://blog.chainalysis.com/reports/cryptocurrency-crime-2020report>). In arguably the most prominent Bitcoin campaign to date, Hamas's military arm, al Qassam Brigades, used its main webpage, social media channels, and officials to call for donations following an Israeli embargo on external financial aid (Katisiri (2019)). Although social media activists helped publicize the campaign, Hamas's main leaked wallets collected a small amount of donations. Al Qassam's unsuccessful campaign suggests that donations may not be the primary use of cryptocurrencies by terrorist groups. Therefore, studies based on leaked addresses likely underestimate the real size of terrorist crypto financing. These activities neither suggest any ability to detect or prevent on-the-ground terror attack operations, nor indicates whether Bitcoin is used to finance on-the-ground attacks.

on-the-ground operations of international terror attacks in the public blockchain system could possibly be because these operations are simply not financed with Bitcoin because this form of financing may put the operation in greater risk compared to the alternatives. In other words, as opposed to donation campaigns for which the transparency of Bitcoin transfers is not likely to risk a specific terror operation, it may be inefficient for terror groups to finance on-the-ground terror attacks using Bitcoin, and they would rather use the traditional banking system for this purpose (Dion-Schwarz, Manheim, and Johnston (2019)). As the old saying goes, public sunlight (transparency) is the best disinfectant. This is also why the existing evidence that bitcoin is used to pay for other crimes such as cybercrime, pornography, or drug purchases (e.g., Foley, Karlsen, and Putnins (2019)) does not imply that on-the-ground terror attack operations are also paid with Bitcoin, as these other crimes usually do not need specialized financing before the action such a large-scale terror event. Using accounting and finance research techniques, forensic accounting, and combining anomaly detection, network analysis, and machine learning on millions of Bitcoin and other cryptocurrency transfers, we fill this gap in the literature.

Large-scale terrorist attacks require financing for several purposes, such as buying weapons and explosives on the black market, buying equipment, and paying the operatives and their families. If financiers of the attack are worried about eliminating transfers trace in the blockchain, they will likely engage in money laundering techniques. These techniques consist of repeatedly reshuffling cryptocurrencies and transferring them across several crypto wallets in the blockchain. For instance, assume that 10,000 dollars are used to finance a terror attack, say to buy machine guns on the black market. If this amount is reshuffled 100 times in several wallets to eliminate its traces, over a million dollars is generated in volume due to the laundering process. We build on this rationale to argue that one would observe abnormally large volumes of transfers in the vicinity of terrorist attacks. Specifically, we expect abnormally large volumes to appear before events occur. Using abnormal volume to investigate an event's information content has been used in the accounting literature dating back to Beaver (1968).

To conduct the empirical analysis, we aggregate millions of transactions at the user level for hundreds of users in the Bitcoin blockchain and classify them into six groups based on their characteristics: exchanges, dark markets, mixers, mining, gambling platforms, and other services. Users in the dark market

group provide services of usually selling illicit products or services on the dark web. Users in the exchange group provide exchange services such as converting fiat to crypto, crypto to fiat, and crypto to crypto currencies. Users in the gambling group provide online casinos and betting platforms. Users in the mining group mine Bitcoin and other cryptocurrencies. Users in the mixer group provide tumbling services, which consist of reshuffling cryptocurrencies into hundreds of transactions and interpolating transactions with other users to decrease or eliminate traceability. Users in the service group provide general Bitcoin services, such as online payments, transfers, and cold storage. We merge these data with terrorist attack data compiled from the comprehensive list of attacks from the Global Terrorism Database (GTD) for 2015-2019.

In our main analysis, we examine Bitcoin's cumulative abnormal volume (CAV) responses in the vicinity of terrorist attacks. We predict and find a sharp increase in CAV responses in the period preceding the attacks, followed by a decrease in the following weeks. Furthermore, we predict and find that the positive CAV response in the days preceding the attack is stronger for events that are likely to require more financing (e.g., greater casualties). We then analyze CAV responses across different users by allocating users into groups of similar services. Mixers are designed to camouflage activity, usually transactions in the dark markets (Foley et al. (2019)). There is also significant variation in the quality of regulation of crypto exchanges and the level of compliance procedures such as KYC and AML (Amiram, Lyandres, and Rabetti (2021)). Thus, we expect abnormal volume to concentrate on crypto exchanges and mixers in the vicinity of attacks. We find that users in the exchange and mixer groups present large positive CAV responses to attacks in the weeks preceding the event, followed by CAV decrease in the week afterward. We also predict and find that the unexpected increased demand for mixer services in the weeks preceding large-scale terrorist attacks leads to higher mixer fees and a slower reshuffling process, crowding out the less time-sensitive volume in the dark markets and other services. Collectively, the results provide evidence of Bitcoin's usage in financing large-scale terrorist attacks.

Since the previous results show that funds likely associated with large-scale terrorist attacks pass through exchanges and exchange-like services, and these businesses are responsible for the largest portion of the transactions during the period, we test the cross-section effects of terrorist attacks by groups (e.g.,

ISIS and al Qaeda) in a partition composed by exchange users. We indeed find that ISIS- and al Qaeda-claimed attacks overseas drive most of our results on exchange users. These terrorist groups have, on average, 21% and 15% cumulative abnormal volume responses in the weeks preceding large-scale attacks, respectively. We do not find economically significant Bitcoin CAV responses to Boko Haram and ELN attacks, used as a placebo test, which are less cryptocurrency literate.

We complement our main analysis with a comprehensive forensic examination of the largest attack in our sample period: the Sri Lanka Easter bombing. This attack, which left several dead and hundreds injured, occurred on April 21, 2019. Besides being the largest ISIS-claimed attack in foreign territories, the attack had a large level of logistics in place (multiple bombing and terrorists on-the-ground), and it has been claimed across several sources that Bitcoin was used to finance the event.⁷ Unlike our main research design, we now identify users with abnormal transfers in the vicinity of the event by applying forensic accounting, and anomaly detection techniques to the user level's aggregated daily transfer time series. We then train machine-learning algorithms with the insights of this analysis and a blockchain based-model. We split the data into training and validation sets and find that, out of the sample, the model predicts terrorist attacks a day before they occur, suggesting that the public information available in the blockchain system is transparent enough to identify patterns associated with large-scale terrorist financing.

Focusing on this suspicious user, we assess whether addresses linked to its wallets have a history of association with other crimes. We track transfers backward in time and find evidence of associations with several reported crimes, including ransom for kidnapped children in Africa and funding jihadi cells in Syria. We also track the funds transferred after the Sri Lanka terror attack and find that some of these funds are likely converted from Bitcoin (BTC) to Ripple (XRP). We trace the funds one step further into the Ripple network and identify a chain of transfers that resemble money laundering, including one anonymous wallet serving as a deposit bank with over 200 million US dollars in reserves. These findings provide further evidence of terrorist crypto financing. We use the insights of this analysis to construct a model predicting terrorist attacks.

⁷See, for instance, <https://en.globes.co.il/en/article-exclusive-isis-funded-sri-lanka-bombings-with-bitcoin-donations-1001284276>.

We implement machine learning algorithms to construct our prediction model. We consider three machine learning algorithms: Supported Vector Machine, Neural Networks, and Random Forest, as they stand out as state-of-the-art solutions for supervised nonlinear learning classifiers. We train these models in the gateway flagged with anomalous transfers around the Sri Lanka Easter bombing. This user provides a set of more than 7 million Bitcoin transfers from 2015 to the end-2019. We split the data into training and validation sets, the former (latter) having about 70% (30%) of the observations for the period. Although all three models have predictive power, the random forest model provides the best performance in the training set. We find that, out of the sample, it predicts terrorist attacks a day before they occur, suggesting that the public information available in the blockchain system is transparent enough to identify patterns associated with large-scale terrorist financing a day before it happens. The model potentially provides tools for governments, agencies, and market participants interested in monitoring, isolating, and shutting down terrorist backers.

The paper makes several contributions to the literature. First, we contribute to the literature on transparency and disclosure choice. As discussed above, the terror financiers essentially tradeoff disclosing more detailed information privately against disclosing less information publicly. The choice of what information to reveal and to whom has been heavily studied in accounting in different contexts (see [Beyer et al. \(2010\)](#) for a literature review). Our study provides evidence that by opting for a blockchain-based currency financing route, the terror financiers leave public evidence of their actions that can be studied, identified, and predicted.

Second, our study contributes to the literature on the interplay of transparency and regulation. An extensive literature in economics, finance, and accounting emphasizes the role of transparency in resolving frictions. For example, this literature suggests that increased transparency can help mitigate systemic risk in the international financial system (e.g., [Bouvard, Chaigneau, and Motta \(2015\)](#)). Relatedly, the role of transparency in the absence of regulation has also been examined. For instance, [Sivakumar and Waymire \(2010\)](#) study voluntary disclosures made by NYSE firms in the early 20th century when accounting standards were not yet in place. [Mahoney \(2009\)](#) questions securities law development and its role in addressing information asymmetries when contract, fraud, property, and company law exists. We add

to this literature by finding that increased transparency in the international fund transfer market allows outsiders to reveal evidence of terror financing and predict terror attacks. Moreover, our study addresses the call for more research on the interplay of transparency and regulation. For instance, [Leuz and Wysocki \(2016\)](#) argue that “*we generally lack evidence on market-wide effects and externalities from regulation, yet such evidence is central to the economic justification of regulation.*” We help to fill this gap by providing evidence that the lack of regulation permits bad actors to exploit blockchain-based currencies in the short-term. However, transparency has the potential to dissuade these users from the system in the long-term as detection risks the operational secrecy of the attack, which is required for operational success. In addition, our paper also answers [Leuz \(2018\)](#) call for more evidence-based policymaking research. Our findings shed important light on the current debate over blockchain-based currency regulation ([Amiram et al. \(2021\)](#), [Cong, Li, Tang, and Yang \(2021\)](#), [Griffin and Shams \(2020\)](#), [Foley et al. \(2019\)](#), [Fusaro and Hougan \(2019\)](#), [Sokolov \(2020\)](#), and [Makarov and Schoar \(2020\)](#)). We do so by indicating that regulation should focus on exchanges and exchange-like services to approximate know-you-customer and anti-money laundering practices to international banking standards.

Third, our study joins the nascent literature examining the role of transparency in alternative financing markets. For instance, [Michels \(2012\)](#) examines voluntary disclosure’s effects in attenuating market inefficiencies in the peer-to-peer lending markets. [Bourveau, De George, Ellahie, and Macciocchi \(2019\)](#) examine crypto analysts’ role in assessing initial coin offering (ICO) projects’ quality. [Cascino, Correia, and Tamayo \(2019\)](#) study the interplay of disclosure and consumer regulation in the reward crowdfunding markets. [Lyandres, Palazzo, and Rabetti \(2021\)](#) study the effects of the disclosed information on crypto-based projects’ operational and financial performance. We contribute to this literature by providing the first analysis of how blockchain transparency can increase scrutiny by permitting outsiders to study the flow of funds and users’ interactions.

Finally, this paper contributes to terrorist financing literature. This literature lacks empirical research, and the reason is apparent: transfers are camouflaged using a vast array of strategies, impeding estimates of the real size of the financing. Therefore, the typical research method is qualitative. For instance, [Rudner \(2010\)](#) surveys the different methods that Hezbollah employs to divert funds, such as charitable

and nongovernmental organizations' contributions. [Pieth \(2002\)](#) discusses whether financial institutions can contribute to the suppression of terrorist financing. Other studies focus on policy considerations. [Schott \(2006\)](#) provides a source of practical information for countries to fight money laundering and terrorist financing. [Navias \(2004\)](#) reviews international, regional, and national efforts to constrain terrorist exploitation of the global financial system before and after Sept. 11, 2001. Our paper makes a novel empirical contribution to this literature.

1 Data

1.1 Blockchain Data Acquisition and Classification

The pseudo-anonymity of Bitcoin transactions imposes nontrivial challenges to identify market participants. First, owners of user wallets must be identified off-blockchain. Second, since each wallet may contain hundreds, if not thousands, of addresses, identifying all addresses is challenging. However, Bitcoin explorers, specialized web pages, social media, and forums have already linked thousands of addresses to popular Bitcoin services, such as exchanges, miners, mixers, dark markets, and services. The methodology of recomposing user wallets from transaction hashes involves two steps. First, to identify a wallet's ownership, we take advantage of the fact that bitcoin is only pseudo-anonymous, which means that all Bitcoin transactions are stored publicly and permanently on the network (anyone can see the balance and transactions of any Bitcoin address); however, the identity of the user behind an address remains unknown until information is revealed. The connection between wallet ownership to a real person or business identity may exist because they are registered in an exchange, leaked to social media or other websites by those who transact with the wallet, or the wallet belongs to business entities that provide their addresses for services (e.g., crypto exchanges). For instance, an online shopper gets access to several addresses from Bitcoin providers (e.g., exchanges), enabling those users' addresses to be revealed. Second, once the addresses are obtained, data fusion algorithms are employed to associate several addresses with one specific user.

For our choice of data fusion methods to aggregate addresses to the user level, we decided to implement

the union-find algorithm.⁸ The main reason is that the algorithm, or similar versions of it, has been used in several academic applications (e.g., Kappos, Yousaf, Maller, and Meiklejohn (2018), Tasca, Hayes, and Liu (2018) and Foley et al. (2019)). Additionally, the union-find algorithm provides a more conservative clustering method, because it is less prone to incorrectly clustering sets of transactions that involve more than one user (Meiklejohn, Pomarole, Jordan, Levchenko, McCoy, Voelker, and Savage (2018), Foley et al. (2019)). This algorithm is used to identify the wallets of 339 Bitcoin users.⁹

A user in our dataset is a collection of wallets for a business entity and may contain thousands of transactions and addresses. For instance, about two hundred thousand addresses and over eight hundred thousand transactions mapped for Binance exchange are treated as one user in our dataset. Since we have the main business players in the Bitcoin ecosystem, we likely cover most of the transactions by known players in the sample period.¹⁰ To classify Bitcoin users into business groups, we follow practitioner and prior research.¹¹ We classify users according to their business type and allocate them into five portfolios: *Exchange*, *Mixer*, *Gambling*, *Dark market*, and *Service*, where *Exchange* contains crypto exchanges responsible for providing exchange and withdrawal services; *Mixer* relates to business associated with reshuffling Bitcoins into several transactions and wallets to eliminate traces in the blockchain; *Gambling* includes gambling, rewarded games, and betting platform; *Dark market* represents online shops, usually on the deep web, responsible for selling illicit goods; and *Service* are general business transactions such as online payments, credit cards, and wallet storage.

⁸The algorithm was designed by Cormen, Leiserson, Rivest, and Stein (2001) and, in the context of cryptocurrencies, first applied by Ron and Shamir (2013). There are several types of data-fusion algorithms and identifying strategies. See <https://en.bitcoin.it/wiki/Privacy>.

⁹Consider the following as a simplified explanation of how the algorithm works. Say the addresses Anne and Bob co-spent (e.g. if the transaction is 1,000 each one sends 500 units) in the transaction Alpha, and the addresses Bob and Carl co-spent in the transaction Beta. Since Bob appears in two different transactions co-spending separately with Anne and Carl, all addresses (Anne, Bob and Carl) must belong to the same user. By repeating this process several times in a pool of millions of transactions, one can aggregate a large portion of addresses to a user data level.

¹⁰For comparison purposes, Foley et al. (2019) exclude users in which transfer of cash does not involve the acquisition of goods or services from its sample. Although they do not mention the number of users, they provide the number of transactions excluded or a total of 88.4 million transactions. We managed to obtain 105.67 million transactions to the same group of users, which means that our sample construction is about 20 percent larger than theirs, even though we cover a smaller period (2014-2019) than them (2009 to 2017)).

¹¹We follow <https://www.walletexplorer.com> classification of users into *Exchange*, *Service*, *Mining* or pools, and *Gambling*, while re-allocating other wallets from the historic group. Additionally, we follow Foley et al. (2019) and add *Dark market* as a category. Assigning users into groups is also present in the accounting literature, for instance, Bushee (1998) assigns investors to dedicated, transient and quasi-indexing investor groups to examine their influence on RD investment behavior.

1.2 Blockchain Data Processing and Statistics

The data collected presents a series of challenges. First, some users become inactive. For instance, Silk Road (the largest dark market responsible for selling illicit goods) became inactive after the FBI seized it. That translates to user sample variation around terror attack events, with some users absent around the most recent attacks. Second, even for active users, several users have no transactions in the vicinity of some events. This usually happens when Bitcoin transactions are infrequent due to lower customer demand or the service accumulating transfers before transferring them to another user. Therefore, from our initially obtained data of 339 users, we apply the following process to construct the final sample used in the paper. First, we exclude a mining wallet that had 613,113 transfers with zero interactions with other users. Second, we exclude 668,768 cold storage addresses, mostly from exchanges and services, that are not relevant for our analysis, as they are used only to store coins. Finally, we also drop transactions that had no interaction with other users (e.g., no output). The total amount of transactions excluded represents only a tiny fraction (0.55%) of the data obtained.¹² The final dataset contains 338 users responsible for engaging in 135.75 million transactions from 2014 to 2019.

Table 1 - Panel A presents the final Blockchain sample summary statistics and each step of the screening process. Users have, on average, 292 thousand unique addresses resulting from data-fusion of 136 million transactions from 2014 to 2019. An extensive set of unique addresses and the total amount of transactions indicate the complexity of mapping the 338 large users in this dataset. *Exchange* and *Service* have the largest number of addresses across groups, being 52 and 32.5 million, respectively. However, when volume is considered, then *Dark market* becomes the second largest group with 31.21 million Bitcoins. *Exchange* leads the amount of Bitcoin transferred with 155.31 million units or over 70% of the transfers in volume during the period. While some transfers contain large amounts of Bitcoins (mean of 10.24), most of the transfers are very small in size (median of 0.01) across all users in the dataset. Notably, users in the *Dark market* group have the largest number of Bitcoins transferred per transaction (mean of 23.7), suggesting that this type of business accumulates a large amount of bitcoin before cashing it out or converting to other cryptocurrencies. *Exchange* users tend to keep a large balance of Bitcoins than users

¹²Our inferences are unchanged if these transactions are not excluded from the final sample.

in different business niches. That is consistent with exchanges needing reserves to operate their services and maintain a minimum amount of liquidity. The average life varies across business types. For instance, *Dark market* has the most extended average life (4.3 years) among all groups. The reason is that these were the first type of businesses to adopt Bitcoin. In contrast, *Mixer*, a service where deposited dirty coins (coins linked to illicit activities or additional privacy) are mixed into several transactions and aggregated back into clean coins (difficult to trace coins), has the shortest average life (2.7 months). The relatively short life is consistent with this type of service reducing Bitcoin's traceability.

1.3 Terrorist Events Data

To create a list of terrorist events, we rely on the Global Terrorism Database (GTD), a database with comprehensive details about terrorist attacks curated at the University of Maryland and used in several academic papers (e.g., Cuculiza, Antoniou, Kumar, and Maligkris (2020)).¹³ The dataset of terrorist attacks contains 7,212 events from 2015 to 2019 that occurred in several countries. We collect information on the number of dead and injured, location, and perpetrator for each event. This information is important to classify terrorist attacks, for example, by terrorist group name or to a certain level of destruction (proxied as the sum of dead and injured). We perform the following filtering process to obtain our final sample. We exclude 3,546 attacks inside Syria, Iraq, Afghanistan, Somalia, Nigeria, and Colombia because these attacks are carried out primarily by local groups. This restriction is necessary because local attacks are less likely to need Bitcoin financing since other means of funding are already in place. We also exclude 3,239 events because the size of destruction, captured as the sum of dead and injured, is smaller than 20.¹⁴ The idea is to keep only large-scale attacks in which financing is more likely to be in place. For the same reason, we exclude 376 attacks that require a low level of logistics (e.g., stabbing or arson) and attacks related to wars, conflicts, and clashes. Finally, we exclude seven attacks in which the date of occurrence overlaps with the event window of other attacks.¹⁵ The last filter allows us to prevent that two or more

¹³For more information, please access <https://www.start.umd.edu/about/about-start>.

¹⁴In later tests, we predict and find the same directional, but economically smaller, effect as our main results for attacks with lower than 20 dead and injured.

¹⁵In this case, we keep attacks that generated a larger number of injured and dead.

similar attacks cluster in time, contaminating our event study analysis. The resulting list of terrorist attacks used in our event study's main specification includes 44 large-scale attacks resulting in an average of 99.12 total injured and dead per attack. [Table 1 - Panel B](#) reports each step of the filtering process. The table also reports the distribution of attacks per size used in a less restrictive analysis. The distribution of terrorist attacks per size follows 54.3%, 28.7%, and 16.9%, for small, mid, and large terrorist events captured as less than 5, between 5 and 20, and above 20 total number of dead and injured. We also filter the list of events for perpetrators when studying the cross-variation effects on claimed attacks. The filtering follows two criteria: terrorist attacks exclusive to the examined terrorist group and the top decile by the destruction level. In this specification, the resulting lists of terrorist attacks by terrorist groups include 85 attacks claimed by ISIS (30), al Qaeda (20), Boko Haram (28), and ELN (7).

2 Predictions to the main analysis

Our goal is to examine whether outsiders can exploit blockchain's transparency and identify flows of funds associated with terrorist attacks. We argue that large-scale attacks need financing for purposes such as buying weapons on the black market and paying terrorists on the ground. If the terror financiers are worried about eliminating traces in the blockchain, they may use money laundering techniques.¹⁶ For instance, assume that 10,000 dollars are used to finance a terror attack, say, to buy a machine gun on the black market. If this amount is reshuffled a hundred times in several wallets to reduce its traceability, over a million dollars are generated in volume due to the laundering process. We build on this rationale to construct the notion that we would then observe abnormally large volumes in the vicinity of terrorist attacks. Therefore, our first prediction is that an abnormally large volume appears before the event happens.

Although we also expect the volume after the attack to return to its normal level, there may be other reasons for it to decline even further. The decline below the normal level could occur because large-scale terrorist attacks are negative news, potentially suppressing Bitcoin trading activities. Alternatively, the

¹⁶An alternative to cryptocurrencies is the Hawala system. It consists of an international network of money brokers operating outside of, or parallel to, traditional banking. This network might be used to facilitate money laundering, avoid taxation, and move wealth anonymously. Following 9-11, the FBI increased pressure to prevent the system from functioning inside the United States. However, al Qaeda used the Hawala system to remove Afghanistan's funds as U.S.-led forces swept through the country in 2001 ([Wheatley \(2005\)](#)).

suppression may reflect a reduction in activity in unrelated illicit transactions that seek to be under the radar from the scrutiny of agencies searching for terrorist linkages due to the attack. Nevertheless, some reasons may cause the volume after the event to remain abnormally high. If the terror financiers pay to the perpetrators' families (e.g., compensation for terrorist death) after the event using bitcoin, we may observe a higher-than-normal volume. To conclude, bitcoin volume patterns in the weeks after a large-scale terrorist attack are potentially confounded with these factors; therefore, we focus our predictions on volume patterns in the weeks before the events happen.

Our second prediction is that abnormal volume is increasing in terror financing needs (proxied by the sum of dead and injured). We expect that small terrorist attacks are unlikely to be financed or are financed with small amounts of Bitcoin. The reason is that these attacks are usually carried out with a low level of equipment and logistics, typically leading to very few injured or dead. Mid and large-scale events, those that lead to a more significant number of injured and dead, are usually the consequence of a more significant level of logistics (e.g., several perpetrators and organized actions) and equipment employed (e.g., mass shooting and bombing); therefore, demanding a more substantial amount of funds to finance the operations.¹⁷

Our third prediction is that Bitcoin transfers associated with terrorist attacks are likely to go through *Exchange* and *Mixer*. While the former provides users with the option to cash out or to use its exchange wallet for payments and transfers, the latter is used to make Bitcoin tracing in the blockchain more difficult. We assume that while terrorist organizations use professional mixers or mixers-like services to reshuffle Bitcoins in an attempt to make it untraceable, terrorists on the ground uses exchanges for payments and withdraws. Since these services are interconnected in the blockchain, we expect to find abnormal volumes concentrated in the vicinity of large-scale terrorist attacks, mainly in identified *Exchange* and *Mixer* wallets.

Our fourth prediction is that services and users dependent on mixer services are expected to be negatively affected in the vicinity of terrorist attacks. During normal times, *Mixer* activities are mainly used

¹⁷We acknowledge that the number of injured and dead is a noisy proxy for the costs of the attack. There could be a costly attack that eventually left very few dead and injured and vice versa. Nevertheless, the noise in our proxy makes it more difficult to find a relation between the proxy for the cost of the attack and abnormal volume.

to camouflage transfers between users and *Dark market* or other *Service* wallets (Foley et al. (2019)).¹⁸ But since *Mixer* volume is expected to increase in the vicinity of terrorist attacks, due to greater demand from terrorist backers, increased fees and slower reshuffling services would then temporarily crowd out the less-time sensitive *Dark market* and *Service* users.¹⁹ Therefore, we predict lower than expected volumes on *Dark market* and *Service* in the period before the attack takes place. Although a higher demand by terrorists crowds out users of *Dark market* and *Service*, we expect no effect on *Gambling* users, as they do not have an immediate need for mixer services. Gambling platforms usually provide user accounts and credits for gambling and playing. This process is generally carried out some time in advance, as users do not often cash out their proceedings due to fees and other costs.

Our fifth prediction is that terrorists experienced in conducting attacks overseas and with a presence in the crypto space, such as ISIS and al Qaeda, are likely to finance their activities with cryptocurrencies. The reasons are the following. First, they have a history of carrying out attacks in locations distant from their home territories, which requires a means to camouflage transfers of funds abroad. Second, they are present in the crypto space by running donation campaigns, which indicates that they are familiar with blockchain technology. Third, sanctions and territorial losses prevent them from financing foreign activities by conventional methods. Unlike these terrorist organizations, Boko Haram and the National Liberation Army (ELN) are known for having terrorist operations inside their home territories and there is no evidence that they use cryptocurrencies.²⁰ While Boko Haram's main actions fall within Nigeria's borders, ELN's battleground is in the areas surrounding large cities in Colombia. As we expect that Bitcoin is most likely used to finance activities overseas, Boko Haram and ELN terrorist attacks should not be associated with our prediction; therefore, they potentially work as placebo tests to our empirical examination.

¹⁸Some exchanges, such as Coinbase, ban users that transfer coins directly to dark market services. Therefore, users tend to use mixer services to eliminate the link.

¹⁹*Mixer* has been shown to be a faulty service for clients when demand increases (de Balthasar T. and Hernandez-Castro (2017)).

²⁰ELN, in Spanish, Ejercito de Liberacion Nacional, or in English, National Liberation Army.

3 Main analysis - Event study

The accounting literature has used abnormal volume as a measure of informativeness in event studies as early as [Beaver \(1968\)](#). Trading volume in stocks has been associated with the magnitude of surprises in annual earnings announcements ([Bamber \(1986\)](#)), stock price anomalies ([Core, Guay, Richardson, and Verdi \(2006\)](#)), the role of media in disseminating news ([Rogers, Skinner, and Zechman \(2016\)](#)) and investors' reaction to blockchain-related disclosures ([Cheng, De Franco, Jiang, and Lin \(2019\)](#)).²¹ To test our predictions, we use an event study that examines Bitcoin volume in the vicinity of large-scale attacks.

We reason that if a terrorist attack is financed with Bitcoin and money laundering techniques are to eliminate traces in the blockchain, this process generates large volumes. One possible technique, for instance, is the use of a mixer or mixer-like service that consists of reshuffling Bitcoin amounts hundreds of times among several addresses in the chain. This technique potentially creates larger volumes around terrorist attacks. This assumption motivates our rationale for the event study. For illustration purposes, see [\(Figure 1\)](#) as an example of mixed transactions for Bitcoin donors to the terrorist group al Qassam Brigades. Each node corresponds to an address in the Bitcoin blockchain, while a cloud of nodes consists of reshuffled transfers associated with the laundering process. Several nodes that ended transferring funds to al Qassam Brigades wallet have funneled donations for dozens or hundreds of other nodes as visually perceived. A total of 920 addresses and 1,939 transfers were used to generate the final 51 transfers that moved 4,247 dollars to al Qassam's wallet. This is a ratio of 38 reshuffled transfers per final transaction, which generates $38 \times 4,247 = 161,386$ (16) in \$ (BTC) volume. In this case, reshuffling is done for a donation campaign with different characteristics than on-the-ground terror attack financing. More precisely, donation campaigns that use smaller amounts over time, are less time-sensitive and less risky in nature, therefore demanding less money laundering (mixing) activity (to make it hard to trace back the origin of funds on the Bitcoin blockchain) than financing terrorist attacks.

Our measure of Bitcoin volume is the daily sum of total inbound and outbound transfers at the user level. We first estimate the expected volume as the mean volume in the 20 days before the first day in

²¹See also [Bamber \(1987\)](#), [Bamber and Cheon \(1995\)](#), and [Lee \(1992\)](#).

the event window. The event window employed in this study is two weeks before (after) the event. We then mean-adjust the realized volume in the event window by subtracting from the realized volume the mean-estimated volume, both in logarithmic terms (Beaver (1968), Copeland (1979), Bamber (1987)).

$$AV_{u,t} = \ln V_{u,t} - \ln \hat{V}_{u,t}$$

where abnormal mean-adjusted volume AV is calculated at the user u and time t . The average abnormal mean-adjusted volume AAV on a given day t is calculated by summing the abnormal volume for each user in the group and dividing by the number of users in the group N . The cumulative abnormal mean-adjusted volume CAV is also constructed by the sum of AAV in the specified event windows T .

$$AAV_{u,t} = 1/N \sum_{t=1}^N AV_{u,t}$$

$$CAV_{u,t} = \sum_{t=1}^T AAV_{u,t}; \quad t = 1, \dots, n.$$

We report the mean CAV for the period before ($t - 15, t - 1$) and after ($t + 1, t + 15$) the event to measure abnormality. We base the choice for the event window on two considerations. First, we consider the focus of our examination. There are two types of cryptocurrency transfers associated with terrorism. The first type is through donation campaigns, which happen in the long term and between donors and the terrorist organization (e.g., the DOJ intervention). The other form is direct financing between the terrorist organization and terrorists on the ground. The latter has a short-term characteristic (e.g., the machine gun purchase a week before the Paris attack). In the lack of substantial evidence on terrorist mechanics to fund terrorists on the ground, the selected event window seems plausible to address our research question. The second consideration is that even when filtering terrorist attacks to at least more than twenty dead and injured, terrorist events are nearly monthly. Therefore, we decided on the event window's length so that the estimated window ($t - 35, t - 16$) is not contaminated with any other similar event.

Our first empirical examination studies the mean CAV responses to 44 large-scale terrorist attacks across all users in our sample. As discussed earlier in the Data section, the vector of terrorist events is filtered for large-scale terrorist attacks that occurred far away from terrorist territories and with a high logistics level. [Table 2 - Panel A](#) reports the results. The mean CAV for the weeks preceding terrorist

attacks is on average 47%. The increase in CAV responses before the event occurs is consistent with the increase in money laundering activities to attend unexpected demand from terrorist financiers. The mean CAV responses fall immediately after events occur, resulting in an average of -43% across all users in our dataset. [Table 2 - Panel B](#) plots the CAV responses' evolution. These results are consistent with our prediction that Bitcoin's abnormal volume is associated with large-scale terrorist attacks.²² As noted in our discussions for mean CAV responses in the period after the event, the decline below normal levels could occur because a major terrorist attack is a negative news, and it may suppress Bitcoin activities. Another possible reason for the further decrease is a reduction in activity in unrelated illicit transactions that seek to be under the radar from the scrutiny of agencies searching for terrorist linkages (e.g., [Foley et al. \(2019\)](#) point out that over half of Bitcoin transactions are linked to illicit activities).²³

Moreover, the results permit us to estimate the average cost of large-scale terrorist attacks. Given our estimated abnormal volume for the event window and the average price of Bitcoin during the period, the estimated upper-bound average cost of funding for large-scale terrorist attacks is 105 thousand dollars. We use the following assumptions to estimate the average cost of a large-scale terrorist attack that occurs far away from terrorist territories and with a high level of logistics in place. First, the average nominal normal volume in the vicinity of the events is 13,000 BTC. Since transaction volume increases by 22.42% in the vicinity of large-scale attacks on average, the estimated abnormal nominal volume is about 2,900 BTC. However, this figure is in volume terms, and we assume that Bitcoin is being reshuffled to make it difficult to trace on the blockchain (that includes but is not necessarily limited to mixers). Deflating the volume by its estimated laundering process results in a cost of about 29 BTC. At an exchange rate of \$3,600

²²A possible concern with this result is that large-scale terrorist attacks occur at a higher rate at the end of the year holidays, since terrorists aim to maximize the number of deaths, when Bitcoin volume may behave differently (e.g., fewer trading activities). Therefore, as a robustness test, we exclude attacks happening in December and January (about 19% of the events). Inferences remain unchanged when we exclude these attacks.

²³To test the robustness of our findings, we estimate the mean CAV responses across different model specifications (Un-tabulated). First, we expand the event window to $(t - 30)(t + 30)$ and to $(t - 45)(t + 45)$. The mean CAV responses for expanded 30 days event windows are virtually the same as the main specification. However, as we expand to 45 days event windows, the mean CAV response for both before and after periods significantly increases. The latter result suggests induced volatility in the event window due to overlapping events, which leads to overestimated coefficients. We also expand the estimated normal windows to 90 and 180 days, respectively. Expanding the estimated normal window, which incorporates other similar events, has a similar effect of inducing volatility in the mean CAV responses; but, it is more pronounced in the period after the events occur. We conclude that the main specification results are robust to both expanded events and estimated normal windows. Yet, due to events clustering in time and the potential contamination of the results on longer windows, the narrow window we use seems adequate.

per Bitcoin during the period, the average large-scale terrorist attack upper bound cost in dollars is about \$105,000. But that is assuming that all the abnormal volume we capture corresponds to terrorist financing, which is likely to be biased upward as prior research shows that there is significant autocorrelation in daily volume and between volume and price (e.g., [Amiram et al. \(2021\)](#)).²⁴

Our second empirical test examines whether the abnormal volume before terror attacks predictably varies with a proxy of the cost of the attack. To do so, we use a less restrictive sample of terrorist events as it considers only the destruction magnitude of the attack, measured as the sum of dead and injured. To construct this test, we partition the list of terrorist events into small and mid-scale attacks (determined by whether the total amount of dead and injured is below five or between five and twenty, respectively). This is in addition to our examination of the 44 large scale attacks (over twenty dead or injured) for which we already reported results in Table 2. Since these lists contain an overwhelming number of attacks, we randomly select 44 events, as in our main large-scale attacks sample, from the list in each bootstrap run and report the mean CAV responses after 50 runs. [Table 3](#) report the results for these tests. The mid-size samples' results are statistically significant, with a mean CAV of 2.36 (-2.47) percent in the weeks before (after) of the events take. As expected, the CAV response for these events is smaller than that of large-scale attacks as mid-scale attacks are likely to require less financing. The bootstrapped coefficients for the small-size scale attacks are nearly zero and significant (not significant) in the period before (after) the event takes place. The difference between mid-scale and small-scale terrorist attacks bootstrapped coefficients is 2.09 (-2.42) percent for the period before (after) the event occurs. The difference between the coefficients for large-scale attacks (reported on [Table 2](#)) and mid-scale bootstrapped terrorist attacks is 44.20% (-39.56%) percent for the period before (after) the event occurs. These results are significant and consistent with our prediction that the abnormal bitcoin volume (as a proxy for terrorist financing) is increasing in the estimated size of destruction.

²⁴Obviously, the estimated cost cannot be taken at its face value. We suggest caution with these estimates for several reasons. First, we do not know the exact size of shuffling activities in reality, and our estimated 100x reshuffling may be larger or lower. Second, the costs of financing terrorists attack likely vary according to the operation complexity, location, perpetrators involved, and type of attack. Third, the calculation considers the mean CAV responses in the weeks preceding the attack and discount the historical correlation between Bitcoin price and volume; however, several other factors (not related to terrorist financing) may also affect the nominal volume. To conclude, this back of envelope calculation illustrates our best effort to deliver an estimated economic impact of terrorist financing.

Our third empirical examination focuses on identifying which type of Bitcoin service is more likely to be associated with these attacks. [Table 4 - Panel A](#) reveals that *Exchange* and *Mixer* exhibit large and significant positive CAV responses in the vicinity of the event of 15% and 25%, respectively. CAV responses decrease in the week following the event, as seen in the post-period's negative coefficients. This pattern is consistent with large volumes of Bitcoin being reshuffled in the period just before a terrorist attack, with a subsequent decline of the transfers' activities ([Table 4 - Panel B](#)). The result supports our prediction that exchanges and mixer services are channels for terrorist funding. Moreover, it suggest that current crypto exchange's compliance procedures (AML and KYC) are not enough to curb the financing of terrorist attacks.²⁵

Consistent with, *Dark market* and *Service* CAV responses are negative in the period preceding the event. The mean CAV responses for *Dark market* and *Service* are practically unresponsive in the week after the event. The patterns we document in these wallets suggest that, during regular times, mixer activities concentrate on camouflaging transfers in the *Dark market* and *Service* wallets. However, in the vicinity of terrorist attacks, more laundering resources are needed, and mixer activities shift to *Exchange* wallets. The increased costs of mixer services and slower reshuffling crowd out less time sensitive users, such as *Dark market* and *Service*. Finally, CAV responses for *Gambling* wallets appear only after the event. We conclude that most of the activity likely associated with reshuffling funds to finance terrorists is attributed to *Exchange* and *Mixer* wallets. The result supports our prediction that *Exchange* and *Mixer* services are channels for terrorist funds. It also supports our prediction that terrorist attack funding crowds out the less time-sensitive *Dark market* and *Service* during the event window.

Our fourth test assesses CAV responses within the *Exchange* group to large-scale terrorist attacks, carried out overseas and partitioned by the three most frequent perpetrators: ISIS, al Qaeda, and Boko Haram. In this analysis, we concentrate on the exchange group for two reasons. Our previous results suggest that funds associated with terrorist attacks pass through exchanges and exchange-like services. Also, exchanges are responsible for the largest amounts of Bitcoin transfers interconnecting with all other

²⁵In untabulated results, we find that, as expected, the CAV response preceding the event is significantly larger for unregulated exchanges than for regulated exchanges. However, since the number of regulated exchanges in the sample is low, we do not elaborate on this prediction and results.

types of services in the blockchain.²⁶ The reason for dropping *Mixer* in this analysis is twofold. First, because they are intermediaries that link users and other services, such as Exchange; therefore, by focusing on *Exchange*, we indirectly observe the effect of *Mixer* activities. Second, because one of our focuses is on showing policymakers where regulation should be draft. Mixers are a shady business in nature (e.g., they already wash Bitcoins for illicit purposes). Very little can be done to shut down their activities (e.g., unknown owners and business location). Unlike mixers, exchanges are usually legal business entities, with headquarters and known owners. Thus, regulation, if addressed to exchanges, is more likely to be enforced.

We expect to find evidence for al Qaeda and ISIS because both groups have a history of terrorist attacks overseas, presence in the crypto space through donation campaigns, and limited access to traditional means of financing due to sanctions and restrictions to capital mobility. However, we do not expect to find evidence for Boko Haram as this group is not active in the crypto space and most of its activities occur in Nigeria and bordering countries. As such, examining Boko Haram serves as a placebo test. As an additional placebo test, we also include ELN because this terrorist group carries out attacks only in Colombia and funds its activities mostly with cash from the drug trade. Similar to Boko Haram, we do not expect to find evidence for ELN attacks. To construct the vectors of terrorist events, we filter events by group's claimed responsibility and allow the top decile deadliest attacks for ELN and Boko Haram, as these groups' attacks result in fewer deaths than al-Qaeda and ISIS. This procedure results in four vectors of 30 (20, 28, and 7) terrorist events claimed by ISIS (al Qaeda, Boko Haram, and ELN).

[Table 5 - Panel A](#) presents the main results. The mean CAV responses for al Qaeda and ISIS exhibit, on average, 15% and 21% in the vicinity of their claimed attacks overseas, respectively. The mean CAV responses for the days before the event are slightly above 5 percent for both groups, consistent with terrorist financing. In contrast, while the mean CAV response for al Qaeda decreases, for ISIS increases in the days after the event takes place. [Table 5 - Panel B](#) plots the CAV responses evolution. After the event occurs, ISIS' mean CAV response is not consistent with our expectation that the volume after the

²⁶Exchanges allow users to exchange fiat-to-crypto, crypto-to-crypto, and crypto-to-fiat at near-zero costs. They also provide clients services, such as withdraws, within-exchange transfers, and online payments, which are essential for financing terrorist operations on the ground.

event is likely to return to its normal levels. This is perhaps due to a special compensation scheme in ISIS that provides insurance payments for the perpetrator's family. However, we acknowledge that we cannot test whether the increase in mean CAV after the ISIS-claimed terrorist events is attributed to this type of compensation or other alternative channels. Conversely, Boko Haram and ELN's results suggest their attacks are unlikely to be associated with Bitcoin transfers. Both groups' mean CAV responses before the event occurs are statistically significant but slightly negative. The mean CAV responses for the placebo groups in the days after the event occurs are statistically significant, but the coefficients are small and not economically significant. The results suggest that groups with experience in attacks overseas and a larger presence in the crypto space, such as ISIS and al Qaeda, are more likely to finance their attacks with cryptocurrencies. Altogether, the results in this section shed additional light on the sponsors of terrorist financing.²⁷

4 Complementary analysis - Anomalous transfers

The event study methodology we used in the previous section has two primary limitations. First, parties interested in researching or tracking and shutting down terrorist associated services need to have more practical tools to pinpoint these funds' origin and destination. Second, the research design does not allow us to pinpoint specific users in the Bitcoin network whose activities are associated with terrorist attacks. We address these limitations by carrying out a comprehensive analysis of anomalous transfers in the Sri Lanka Easter bombing's vicinity. This analysis also provides further insights to construct our model for terrorist attack prediction. Our motivation to examine this attack stems from rumors that arose after the event claiming that Bitcoin had been used to finance it.²⁸ However, there is no information about the users and wallets involved in this attack, so we employ several techniques to pinpoint anomalous transfers at the individual-user level and to track the funds onto the blockchain.

²⁷We carried out several other untabulated placebo tests, such as using holiday dates, random dates, or main sport events dates. All tests support our main results in this section.

²⁸See, for instance, <https://en.globes.co.il/en/article-exclusive-isis-funded-sri-lanka-bombings-with-bitcoin-donations-1001284276>.

We follow [Laptev, Amizadeh, and Flint \(2015\)](#) and [Dai, Byrnes, Liu, and Vasarhelyi \(2019\)](#) and employ a rolling three-sigma rule to detect anomalous transfers in individual users in the vicinity of the Sri Lanka Easter bombing. We consider a transfer anomalous if it falls more than three standard deviations from its user's last three months' historical mean. In large samples, a small number of outliers is expected, but, in a time-series, the timing of the appearance of outliers may indicate an anomaly. Outliers, being the most extreme observations, may be suspicious if they appear around specific events, in our case, large-scale terrorist attacks. The rule fits well for our purposes, because the transfer data tend to be normal and large-scale terrorist attacks are rare.

[Table 6](#) reports the results for the three-sigma rule test. The rule identifies 48 users with at least one anomaly during the month of April. Users have on average only 2 percent of the transfers outside three standard deviations. *Dark market* wallets have more anomalies than the other groups on average. The reason may be business-driven, as this type of business probably accumulates a certain quantity of payments before converting Bitcoins to fiat. For instance, some users have a very low frequency of transfers during a certain period but a few spikes in other periods. These spikes mostly relate to transfers between *Dark market* and *Exchange* wallets. *Mixer* anomalies are the largest, followed only by *Exchange*. Results in these groups are driven by a large transfer at the beginning of April, discussed earlier. Most of the anomalies detected in April (38 users) are consistent with the large transfer across US-based businesses at the beginning of the month. Another four users had anomalies detected too far (-15,+15) from the event. The remaining six relevant suspicious users are distributed into *Service* (3), *Exchange* (2), and *Dark market* (1). The largest anomaly in the period, almost 400 BTC (2.8M USD), relates to a popular gateway.

The gateway's main business consists of offering online sellers the possibility to accept cryptocurrencies for payment. For instance, it provides online shopping carts where clients click to buy a product with Bitcoin. The gateway receives the payments in crypto, takes a fee, and sends fiat payments to businesses. By doing so, it eliminates risks of online shopping attributed to oscillations of cryptocurrency prices. The difficulty of accepting Bitcoin as payment is severe, as its volatility is much higher than widely used currencies ([Yermack \(2015\)](#)).

The gateway surfaces with suspicious transfers occurred off three standard deviations from its historical mean in the vicinity of the event. Moreover, these transfers are user-specific; that is, they do not appear in other users in our dataset. Assessing the wallet's balance, we find that the mean balance increases significantly a day before the event, followed by a sharp decrease in balance on the day afterward, which is a pattern consistent with terrorist financing. A series of transactions caused a slow increase and normalization to the mean balance before the event, but a sharp decline is flagged at the event end-day at a value of 400BTC (2.8M USD).

We further investigate whether wallets associated with money laundering techniques are used to feed the gateway's abnormally large transfers. We find a wallet constantly feeding transfers into the gateway. The wallet has over 1 million transactions and approximately 1.9 million worth in Bitcoin inbound and outbound transfers. The difference between inbound and outbound transfers is just less than 0.001 BTC. The extremely high volume of transfers and the fact that inbound and outbound transfers match almost perfectly indicates that the associated wallet is a mixer used to reshuffle funds and is likely related to illicit activities. We also found that this wallet is associated with at least 29 reported crimes, including ransoms for kidnapped children in Africa and the funding of jihadi cells in Syria.²⁹

We continue by exploiting the fact that some of the funds were transferred to a crypto exchange. Three destinations for these funds are likely: (i) exchanged to fiat and withdrawn, (ii) sent to several other addresses, or (iii) exchanged to another cryptocurrency. Unfortunately, testing the first two hypotheses is impractical for two reasons. First, this exchange is one of the largest, containing thousands of daily transactions.³⁰ Second, exchange transfers may occur off the blockchain in the exchange platform. However, the conversion hypothesis can be tested by an association of user, date, and value in other blockchain platforms. To do so, we searched for the gateway's associated wallets in other cryptocurrencies, such as Litecoin, Ripple, and BitcoinCash. A few wallets denominated in these cryptocurrencies are found. However, a wallet in Ripple surfaces as the best candidate because the timing, amount, and direction of the

²⁹An address can be verified for reported crimes, such as extortion and ransomware, in services such as <https://bitcoinwhoswho.com/> and <https://www.bitcoinabuse.com/>. All addresses used in this study are available upon request.

³⁰Although Bitcoin transactions are traceable, Bitcoin units are not. For instance, if two Bitcoins are transferred from wallet A to wallet B containing one Bitcoin and then from wallet B to wallet C, two Bitcoins are transferred. One can be assured that one bitcoin in C came from A but not certain which of the two Bitcoins was originally spent.

transfer match the information we have on the exchange. The account was activated in December 2014 and has over 250,000 transactions and a balance of over 17 million XRP (5.15M USD) at the time of writing.

Table 7 - Panel A shows the daily volume of transactions at the gateway Ripple wallet (GRW) during the first half of 2019. The wallet moved, on average, 133.34 (153.31) thousands XRP in daily inbound (outbound) transfers. However, April 23 is an atypical day in the period, as seen in Figure 2. During common days, GRW presents an average of 200,000 XRP in total transfers. However, most of these transfers are composed of less than 1,000 XRP. A large inbound transfer of 660,000 XRP (212,000 USD) arrived on the morning of that day, stayed at GWR wallet for about half an hour, and then was transferred to an anonymous wallet. The transfer is suspicious for three reasons. First, the Ripple network's inbound amount and timing match the outbound amount and timing in the Bitcoin blockchain. Second, the XRP transfer comes from the same exchange where BTC funds were sent. Finally, these transfers are outliers to the average historical transfer size in the GRW. To understand how abnormal these transfers are, as seen in Table 7 - Panel B, most of the GRW transfers fall below 1,000 XRP. Small transfers (<1,000 XRP) are consistent with the gateway's main business, based on capturing cryptocurrency payments from online stores. Larger transfers may relate to vault storage or other services and are not obviously odd, but extremely large transfers (>100,000 XRP) are inconsistent with the gateway business model.

Additionally, these transfers' nature allows us to track them a step further on the Ripple network. First, practically the same amount that entered the GRW exited it. Second, funds remained in this wallet for less than half an hour with no other transactions in between. The three wallets following the GRW transfer have very peculiar behavior that resembles money laundering. The first wallet behaves as a mixer, reshuffling large sums of funds among several hundred wallets. The second receives reshuffled funds from the mixer and distributes them to several anonymous wallets. Finally, the third wallet behaves as a deposit bank connected to this money laundering chain and has far more inbound than outbound transfers during its lifetime. As of November 2020, the third wallet has a balance of 786M XRP (211M USD), which corresponds to one of the wealthiest wallets on the Ripple blockchain.³¹

³¹The wallet is among the top 0.1% richest wallets on Ripple blockchain <https://ledger.exposed/rich-stats#percentage>. ISIS is claimed to have a missing war chest of about 300 million dollars (see <https://>

To conclude, the evidence in this section suggests that a chain of sophisticated money laundering wallets, in both Bitcoin and Ripple blockchains, participated in the financing of the Sri Lanka Easter bombing. Since the data behavior, such as the type, amount, and timing of transfers and wallets associated with the terrorist attack event, are potentially rich sources for predicting terrorist attacks, we explore this angle in the next section.

5 Additional analysis - Predicting terrorist attacks

In this section, we test whether Bitcoin transfer patterns can predict terrorist attacks. Specifically, we test whether we can use the previous sections' insights to construct a model that predicts terrorist attacks with accuracy better than the historical rate. To do so, we employ three different machine learning models in the training set and choose the best performer to extend the analysis in the validation set. We consider Supported Vector Machine (SVM), Neural Networks (NN), and Random Forest (RF) machine learning models, as they stand out as state-of-the-art solutions for supervised nonlinear learning classifiers. While we must calibrate SVM to provide reasonable results, NN and RF are specification-free. A drawback of NN is that the parameters are harder to interpret than from SVM or RF. Despite these differences between models, we have no prediction of which model will be the best fit for our purposes.

We identify the best model by using similar tuning (model calibration) across these models and choosing the one that generates the best performance. There are several performance metrics that can be used, which we elaborate on below.³² We train these models in the flagged user as we learned, from the previous section, that this user has several anomalous transfers in the vicinity of the Sri Lanka Easter bombing. To define the parameters included in the classification model, we draw predictors from the blockchain data, as we learned from the previous sections that they are associated with terrorist attacks.

[//www.thenationalnews.com/world/fears-missing-isis-millions-are-hidden-in-cryptocurrency-ready-for-use-as-war-chest-1.1021275](http://www.thenationalnews.com/world/fears-missing-isis-millions-are-hidden-in-cryptocurrency-ready-for-use-as-war-chest-1.1021275)).

³²For the SVM model, we adjusted with the Gaussian kernel function to approximate it to NN and RF performances. As training classifiers in high dimensions (several parameters or resampling) is time-consuming, the adjustment helps this model provide faster and more accurate results.

For each model, we run the following classification formula:

$$Error_{(t)} = Volume_{(t-1)} + In_{(t-1)} + Out_{(t-1)} + Life_{(t-1)} + Anonymous_{(t-1)} + Exchange_{(t-1)} + Mixer_{(t-1)} + DarkMarkets_{(t-1)} + Balance_{(t-1)} + Sigma_{(t-1)}$$

Where *Error* is an indicator variable that captures whether a terror attack occurred on that day. The list of terrorist attacks is filtered for attacks with more than 20 dead or injured, claimed by either ISIS or al Qaeda, and that occurred in countries far away from these groups' territories. *Volume* is the logarithm of the sum of inbound and outbound transfers measured in Bitcoin units. *In* is the logarithm of the inbound transfer value measured in Bitcoin units. *Out* is the logarithm of the outbound transfer value measured in Bitcoin units. *Life* is the life of the wallet, defined as the distance in days from the first to the last trade, associated with sending or receiving the funds from users in our sample. *Anonymous* is an indicator variable that captures whether funds came from, or were sent to, an anonymous address. *Exchange* is an indicator variable that captures whether funds came from, or were sent to, an exchange address. *Mixer* is an indicator variable that captures whether funds came from, or were sent to, a mixer address. *DarkMarkets* is an indicator variable that captures whether funds came from, or were sent to, a dark markets address. *Balance* captures the total balance of the user's wallet in Bitcoin units. *Sigma* is an indicator variable that captures whether transfer size falls off three standard deviations the user's last three months' historical mean. All predictors used in the classification model are at the transfer-user level and lagged by one day.

We train this model across the three machine learning algorithms in the gateway flagged with anomalous transfers at the Sri Lanka Easter bombing vicinity. This user provides a set of 7,443,285 transfers in the period of 2015 to 2019. We split the set into training and validation sets, with the former (later) having about 70% (30%) of the observations for the period. That results in 21 (9) events in the training (validation) sample. Rather than manually setting parameters that we believe have higher predictive power, we tenfold cross-validate all three models (James, Witten, Hastie, and Tibshirani (2017)). The cross-validation combines averages of measures of fitness in prediction to derive a more accurate estimate of model prediction performance. Each round of cross-validation involves partitioning the training data into 19 equally sized subsets, performing the analysis on that subset, and comparing it to results from the previous rounds to finally arrive at the model's best fit. Once the final tuning values are assigned, the final

model is refit using the entire training set. The technique also helps to mitigate over-fitting (Kuhn and Johnson (2018)). This process is repeated 100 times for each model.³³

Another concern arises from having too many binary predictors. Since we tune models using resampling methods, a random sample of the training set may result in some binary predictors becoming a zero variance predictor. The near zero-variance predictors can cause numerical problems during resampling for some linear models (Zorn (2005)). We address the issue by implementing nonlinear classification models that are less prone to having numerical issues derived from near or zero-variance predictors. Additionally, we test all binary predictors on whether the percentage of unique values is less than 20% and whether the ratio of the most-frequent to the second most-frequent value is greater than 20. The SVM model best fit is specified with a parameter cost of 0.5 and 3,856 supported vectors, resulting in an accuracy of 0.994 and a Kappa of 0.107. The NN's model best fit is specified with size 5 and decay 0.1, resulting in an accuracy of 0.994 and a Kappa near zero. The RF's model best fit is specified with 500 trees and nine variables at each split resulting in perfect accuracy and a Kappa of 0.99.

Table 8 reports results across all models for the subsample of the training set in which these models have their respective highest performance. The *No Information Rate* (or a naive guess) is a guess based on the historical expected probability that tomorrow has no terror attack. As observed before, the high rate indicates that terror attacks are rare. However, all three models' accuracy suggests that the machine learning models predict better than a without-learning guess. In other words, even for the highly skewed no-terror attack occurrences, these models can pick attacks with accuracy better than a guess without learning. However, as indicated by lower kappa, the SVM and NN results are not significant at 95% confidence level. That is important, because Kappa is a more relevant performance metric than accuracy when classes are highly unbalanced (Landis and Koch (1977)). The confusion matrix, which is the tabulation of model prediction and real outcomes, demonstrates that SVM and NN produced false positives (NO/YES), while the prediction (YES/NO) indicates that both models produced a large number of false negatives. We follow McNemar (1947) to test whether there is a significant sampling error affecting the differences between correlated proportions in the confusion matrix table. We reject the null in all three models. A more

³³Although very unlikely, it is possible to implement resampling incorrectly. A hundred repetitions in the resampling technique mitigate this issue.

appropriate measure is *Balance Accuracy*, which equally weights the accuracy for predicting positive and negative events. For instance, the naive predictor is that there are never terrorist attacks has a balanced accuracy of 0.497 $((0.994+0)/2)$. All models have a better balance accuracy than a guess without learning. The tuning results suggest that the RF model stands out as the best fit for predicting terrorist attacks in the validation set.

We therefore apply our RF-trained model to the prediction set at the gateway and to a group of users defined by their business type: *Exchange*, *Gambling*, and *Service*.³⁴ The model accuracy in these validation sets increases from a flagged user to a group of users (Table 9). The mean sensitivity (precision) is higher for grouped data, which means that the performance is better in detecting terrorist attacks when they indeed happen. However, mean specificity (recall) drops considerably, translating to a larger number of false alarms. Yet, as mentioned before, because the predicted variables' class is highly skewed toward no terrorist attacks, accuracy is not the most appropriate measure. Therefore, we focus on balanced accuracy, as this measure equally weights the accuracy of predicting terrorist attacks and predicting no terrorist attacks. As stated, a naive guess that there is no terror attack is accurate 99.37% of the time, but the balance accuracy for this guess is 49.68%. In this measure of performance, the model has better performance at the flagged user and in the *Gambling* group, is not significantly different than a naive guess in the *Exchange* group, and it has worse performance in the *Service* group. Most of the anomalous transfers flagged by the model (about 370 transfers) correspond to the Sri Lanka Eastern bombing; however, we successfully predict several other events in the period.

Overall, this section shows that anomalous transfers of Bitcoin in the vicinity of terrorist attacks have predictive power. The machine learning models predict whether a terror attack will happen in the next day. The performance is persistent across users and time, indicating that the association between terrorist attacks and these transfers is not spurious.

³⁴We omitted *Mixer*, *Mining* and *Dark market*, because reinforcement is not effective in this type of business; there are too few users, and missing data around selected events affects model performance.

6 Conclusion

In recent years, there has been a significant growth in international money transfers through blockchain-based cryptocurrencies, which essentially use public accounting ledgers. The proliferation of cryptocurrencies as a means to transfer value across borders has two opposing effects related to the transparency of the international financial system. On the one hand, cryptocurrencies and, in particular, Bitcoin may enable criminals to circumvent efforts by governments, financial institutions, and financial regulators to curtail illegal financing. On the other hand, fund transfers that used to be the known only by the parties involved are now transparent on the public blockchain ledger. In this study, we examine whether researchers can identify and predict illegal financial activities, specifically international terrorist attacks, via public blockchain system, by using empirical strategies based on the accounting and finance literatures, forensic accounting, and machine learning algorithms. Large-scale attacks need financing. If an attack's financiers are worried about eliminating transfers' traces in the Blockchain, they likely employ money laundering techniques. These techniques consist of reshuffling cryptocurrencies several times and moving currencies across multiple crypto wallets in the blockchain. We build on knowledge and surmise that researchers could observe abnormally large volumes of transfers in the vicinity of terrorist attacks. Specifically, we expect abnormally large volumes to appear before the event occurs. Using abnormal volume to investigate an event's information content has been used in the accounting literature as early as [Beaver \(1968\)](#).

Our main analysis examines Bitcoin's cumulative abnormal volume (CAV) responses in the vicinity of terrorist attacks. We predict and find a strong positive CAV response in the days preceding attacks, followed by a decrease in the following weeks. Furthermore, we predict and find that the positive CAV response in the days preceding the attack is stronger for events that are likely to require more financing. We then analyze CAV responses across different users by allocating these users into groups of similar services. We predict and find that users in the exchange and mixer groups have large CAV responses to attacks in the weeks preceding the event, followed by CAV decrease in the week afterward. We also predict and find that the unexpected increased demand for mixer services in the weeks preceding large-scale terrorist attacks, which leads to higher mixer fees and slower reshuffling process, crowds out the less time-sensitive volume

in the dark markets and other services. Moreover, we predict and find that ISIS and al Qaeda's overseas attacks drive most of our results. Conversely, as predicted, the results for terrorist attacks claimed by Boko Haram and ELN suggest that their attacks are not associated with Bitcoin transfers. Taken together, this evidence is consistent with cryptocurrencies serving as a means of terrorist financing.

We complement our main analysis by examining the largest attack in our sample period, the Sri Lanka Easter bombing on April 21, 2019. We flag abnormal transactions using forensic accounting and anomaly detection techniques to the user level's aggregated daily transfer time series. Next, we implement machine learning algorithms to construct a terrorist attack prediction model. We train these models in the gateway flagged with anomalous transfers at the vicinity of the Sri Lanka bombing. Although all models have predictive power, the random forest model provides the best performance in the training set. We find that, out of the sample, the model predicts terrorist attacks a day before they occur with fair accuracy, suggesting that the public information available in the blockchain system is transparent enough to identify illegality in real-time. The model also provides tools for governments, agencies, and other market participants interested in monitoring, isolating, and shutting down terrorist backers.

The paper makes several contributions to the literature. First, we contribute to the literature on transparency and disclosure choice. The terror financiers essentially tradeoff disclosing more detailed information privately against disclosing less information publicly. The choice of what information to reveal, and to whom, has been heavily studied in accounting in different contexts. Our study provides evidence that by opting for a blockchain-based currency financing route, the terror financiers leave public evidence of their actions that can be studied, identified, and predicted. Second, our study contributes to the literature on the interplay of transparency and regulation. There is a long literature in economics, finance, and accounting that emphasizes the role of transparency in resolving frictions. Relatedly, the role of transparency in the absence of regulation has also been examined. We add to this literature by providing that increased transparency in the international fund transfer market allows outsiders to reveal evidence of terror financing and predict terror attacks. Moreover, our study addresses the call for more research on the interplay of transparency and regulation. We help to fill this gap by providing evidence that the lack of regulation permits bad actors to exploit blockchain-based currencies in the short term. However, transparency has the

potential to scare off these users from the system in the long term as detection risks the operational secrecy of the attack, which is required for its operational success. Our findings also shed important light on the current debate over blockchain-based currency regulation. We do so by indicating that regulation should focus on exchanges and exchanges-like services to approximate know-you-customer and anti-money laundering practices to international banking standards. Third, our study adds to the literature that examines the role of transparency in alternative financing markets. We contribute to this literature by providing the first analysis of how blockchain transparency can increase scrutiny by permitting outsiders to study the flow of funds and users' interactions. Finally, this paper contributes to terrorist financing literature. This literature lacks empirical research, and the reason is apparent: transfers are camouflaged using a vast array of strategies, impeding estimates of the real size of the financing. Therefore, the usual research method is qualitative. We provide novel empirical contribution to this literature.

References

- Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque country. *American Economic Review* 93(1), 113–132.
- Amiram, D., E. Lyandres, and D. Rabetti (2021). Competition and product quality: Fake trading in crypto exchanges. Working paper.
- Arin, K. P., D. Ciferri, and N. Spagnolo (2008). The price of terror: The effects of terrorism on stock market returns and volatility. *Economics Letters* 51(3), 164–167.
- Bamber, L. S. (1986). The information content of annual earnings announcements: A trading volume approach. *Journal of Accounting Research* 24(1), 40–56.
- Bamber, L. S. (1987). Unexpected earnings, firm size, and trading volume around quarterly earnings announcements. *The Accounting Review* 62(3), 510–532.
- Bamber, L. S. and Y. S. Cheon (1995). Differential price and volume reactions to accounting earnings announcements. *The Accounting Review* 70(3), 417–441.
- Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of Accounting Research* 6, 67–92.
- Belasco, A., M. Eaglen, L. Hartig, T. Jonas, M. McCord, and J. Mueller (2018). Counterterrorism spending: Protecting America while promoting efficiencies and accountability. https://www.stimson.org/wp-content/files/file-attachments/CT_Spending_Report_0.pdf.
- Beyer, A., D. A. Cohen, T. Z. Lys, and B. R. Walther (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics* 50(2-3), 296–343.
- Blomberg, S. B., G. D. Hess, and A. Orphanides (2003). The macroeconomic consequences of terrorism. *American Economic Review* 51(5), 1007–1032.
- Bourveau, T., E. De George, A. Ellahie, and D. Macciocchi (2019). Information intermediaries in the crypto-tokens market. Working paper.
- Bouvard, M., P. Chaigneau, and A. Motta (2015). Transparency in the financial system: Rollover risk and crises. *The Journal of Finance* 70(4), 1805–1837.
- Bushee, B. J. (1998). The influence of institutional investors on myopic R&D investment behavior. *The Accounting Review* 73(3), 305–333.
- Cascino, S., M. Correia, and A. Tamayo (2019). Does consumer protection enhance disclosure credibility in reward crowdfunding? *Journal of Accounting Research* 57, 1247–1302.
- Cheng, S. F., G. De Franco, H. Jiang, and P. Lin (2019). Riding the blockchain mania: Public firms' speculative 8-k disclosures. *Management Science* 65(12), 5901–5913.
- Cong, L. W., X. Li, K. Tang, and Y. Yang (2021). Crypto wash trading. Working paper.

- Copeland, T. E. (1979). Liquidity changes following stock splits. *The Journal of Finance* 34(1), 115–141.
- Core, J. E., W. R. Guay, S. A. Richardson, and R. S. Verdi (2006). Stock market anomalies: what can we learn from repurchases and insider trading? *Review of Accounting Studies* 11(1), 49–70.
- Cormen, T. H., C. E. Leiserson, R. L. Rivest, and C. Stein (2001). Introduction to algorithms. *Cambridge* (MIT Press).
- Cuculiza, C., C. Antoniou, A. Kumar, and A. Maligkris (2020). Terrorist attacks, analyst sentiment, and earnings forecasts. *Management Science* (Forthcoming).
- Dai, J., P. Byrnes, Q. Liu, and M. Vasarhelyi (2019). Audit analytics: A field study of credit card after-sale service problem detection at a major bank. *Rutgers Studies in Accounting Analytics. Emerald Publishing Limited*, 17–33.
- de Balthasar T. and J. Hernandez-Castro (2017). An analysis of bitcoin laundry services. *Lipmaa H., Mitrokovska A., Matulevicius R. (eds) Secure IT Systems. NordSec 2017. Lecture Notes in Computer Science 10674*, Springer, Cham.
- Dhaliwal, D. S., I. K. Khurana, and R. Pereira (2011). Firm disclosure policy and the choice between private and public debt. *Contemporary Accounting Research* 28(1), 293–330.
- Dion-Schwarz, C., D. Manheim, and P. B. Johnston (2019). Terrorist use of cryptocurrencies. *RAND National Security Research Division* (rand.org).
- Eckstein, Z. and D. Tsiddon (2004). Macroeconomic consequences of terror: Theory and the case of Israel. *Journal of Monetary Economics* 51(5), 971–1002.
- Foley, S., J. R. Karlsen, and T. J. Putnins (2019). Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies* 32(5), 1798–1853.
- Fusaro, T. and M. Hougan (2019). Presentation to the US Securities and Exchange Commission. *Bitwise Asset Management*.
- Griffin, J. M. and A. Shams (2020). Is bitcoin really un-tethered? *Journal of Finance* 74(4), 1913–1964.
- ICAEW (2018). Blockchain and the future of accountancy. *ICAEW Thought Leadership - IT Faculty*. (Institute of Chartered Accountants in England and Wales).
- James, G., D. Witten, T. Hastie, and R. Tibshirani (2017). An introduction to statistical learning: with applications in R. *Springer* (7th ed.).
- Kang, C., C. Lee, K. Ko, J. Woo, and J. Hong (2020). De-anonymization of the bitcoin network using address clustering. *Blockchain and Trustworthy Systems. BlockSys 2020. Communications in Computer and Information Science. Springer 1267*, 489–501.
- Kappos, G., H. Yousaf, M. Maller, and S. Meiklejohn (2018). An empirical analysis of anonymity in Zcash. *The 27th USENIX Security Symposium*.
- Karolyi, G. A. and R. Martell (2011). Terrorism and the stock market. *International Review of Applied Financial Issues and Economics* 2(2), 285–315.

- Katisiri, R. (2019). Hamas raises bitcoin donations via US crypto exchange. <https://thenextweb.com/hardfork/2019/04/26/hamas-bitcoin-donations-address/>.
- Kuhn, M. and K. Johnson (2018). Applied predictive modeling. *Springer* (2nd ed.).
- Landis, J. R. and G. G. Koch (1977). The measurement of observer agreement for categorical data. *Journal of Accounting and Economics* 1(33), 159–174.
- Laptev, N., S. Amizadeh, and I. Flint (2015). Generic and scalable framework for automated time-series anomaly detection. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Association for Computing Machinery*, 1939–1947.
- Lee, C. M. C. (1992). Earnings news and small traders: An intraday analysis. *Journal of Accounting and Economics* 15(2-3), 265–302.
- Leuz, C. (2018). Evidence-based policymaking: Promise, challenges, and opportunities for accounting and financial markets research. *Accounting and Business Research* 48, 582–608.
- Leuz, C. and P. D. Wysocki (2016). The economics of disclosure and financial reporting regulation: Evidence and suggestions for future research. *Journal of Accounting Research* 54, 525–622.
- Lyandres, E., B. Palazzo, and D. Rabetti (2021). ICO success and post-ICO performance. Working paper.
- Mahoney, P. (2009). The development of securities law in the united states. *Journal of Accounting Research* 47, 325–347.
- Makarov, I. and A. Schoar (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics* 135(2), 293–319.
- McNemar, Q. (1947). Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* 12(2), 153–157.
- Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage (2018). A fistful of bitcoins: Characterizing payments among men with no names. *Communications of the ACM* 59(4), 86–93.
- Michels, J. (2012). Do unverifiable disclosures matter? evidence from peer-to-peer lending. *The Accounting Review* 87(4), 1385–1413.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash. White paper.
- Navias, M. S. (2004). Finance warfare as a response to international terrorism. *War on Terrorism*, by O’Day, Alan, Routledge. Book Chapter 23.
- Pieth, M. (2002). Financing of terrorism: Following the money. *Springer, Dordrecht*.
- Rogers, J. L., D. J. Skinner, and S. L. C. Zechman (2016). The role of the media in disseminating insider-trading news. *Review of Accounting Studies* 21(3), 711–739.
- Ron, D. and A. Shamir (2013). Quantitative analysis of the full bitcoin transaction graph. *In 17th Financial Cryptography and Data Security International Conference*.

- Rudner, M. (2010). Hizbullah terrorism finance: Fund-raising and money-laundering. *Journal Studies in Conflicts and Terrorism* 33(8), 700–715.
- Schott, P. A. (2006). Reference guide to anti-money laundering and combating the financing of terrorism, second edition. *World Bank*.
- Sivakumar, K. and G. Waymire (2010). Voluntary interim disclosure by early 20th century nyse industrials. *Contemporary Accounting Research* 10, 673–698.
- Sokolov, K. (2020). Ransomware activity and blockchain congestion. *Journal of Financial Economics* (Forthcoming).
- Tasca, P., A. Hayes, and S. Liu (2018). The evolution of the bitcoin economy: extracting and analyzing the network of payment relationships. *The Journal of Risk Finance* 19(2), 94–126.
- Wheatley, J. (2005). Ancient banking, modern crimes: How hawala secretly transfers the finances of criminals and thwarts existing laws. *Journal of International Law* 26, 347–378.
- Yermack, D. (2015). Is bitcoin a real currency? An economic appraisal. *Handbook of Digital Currency*, Chapter 2, 31–43.
- Zorn, C. (2005). A solution to separation in binary response models. *Journal of Political Analysis* 13(2), 157–170.

Figure 1. Al Qassam Bitcoin donation campaign The figure plots the network of transactions of one of al Qassam Bitcoin addresses. Al Qassam address is represented by the red dot at the center of the plot. The blue dots near the center of gravity represents direct donations to al Qassam wallet. The blue dots further away from the center of gravity, represents addresses involved on reshuffling Bitcoin funds. Al Qassam address received 51 final transfers with a total of \$4,247.26 to as August/2019. This plot was generated with 920 Nodes (addresses) and 1,939 Edges (transfers).




Figure 2. Historical Ripple transfers: The plot shows aggregated daily inbound (blue), outbound (red) and total (black) transfers on Ripple network at the Gateway wallet in 2019. XRP transfers are measured in units. An abnormally large transfer (approximately 1.6 million XRP) occurred at the early morning one day after the Sri Lanka Easter bombing.




Table 1. Sample construction: This table reports blockchain sample construction. Panel A reports the number of users, transactions, addresses, and volume for the obtained, excluded, and used blockchain samples. All variables are reported in units, except Transactions, Address, and Volume that are reported in millions. Panel B reports the total, excluded, and final sample of terrorist events. The distribution of events by size follows the sum of total dead and injured in the following way: High-size for more than 20, Mid-size between 20 and 5, and Low-size below 5.

Panel A - Blockchain								
	Users	Transactions	Address	Volume	Average	Median	Balance	Life
Obtained:	339	136.50	99.79	222.24	21.90	0.010	1,429.57	1287
Excluded:								
- mining	1	613.11	668.71	3.32	78.07	15.205	0.06	2220
- cold storages	-	134.86	169.58	1.12	658.03	1.742	83.86	979
- lower interaction	-	0.01	0.01	0.01	182.03	45.106	0.00	125
Used:								
Dark Markets	96	11.26	7.27	31.21	23.69	0.010	10.30	1583
Exchange	97	62.12	51.99	155.31	7.85	0.016	4,884.85	1496
Gambling	49	18.53	6.64	7.78	0.78	0.002	2.83	1345
Mixer	35	0.34	0.28	0.37	6.13	0.010	1.55	81
Service	55	43.50	32.48	23.12	2.00	0.010	165.61	1287
Total	338	135.75	98.95	217.80	10.24	0.010	1,458.19	1290

Panel B - Terrorist Events					
	Excluded	Sample	Injured	Dead	Total
Total number of terrorist events:		7,212	9.55	6.47	16.01
Distribution by size:					
Large-scale event	5,985	1,221	43.56	26.15	69.71
Mid-scale event	5,140	2,072	5.57	4.86	10.43
Small-scale event	3,293	3,919	1.02	1.17	2.19
Restrictions for the main specification:					
Home country	3,546	3,666	7.64	4.24	11.88
Small-scale attacks	3,239	427	47.94	22.30	70.24
Low logistics, war related or arson	376	51	75.92	21.48	97.40
Overlapping	7	44	77.13	23.84	99.12

Table 2. Event study: all users This table reports the cumulative mean-adjusted abnormal volume (CAV) for the periods before, after and whole window (Panel A), and plots the evolution of CAV responses (Panel B), across all Bitcoin services (users). The period before and after is determined in the interval [-15,-1] and [1, 15], respectively. Volume is calculated as the logarithmic change of the sum of inbound and outbound volumes. Abnormal volume is mean adjusted for estimates generated in the period of 20 days before the first day in the event window. The mean-adjusted cumulative abnormal volume is reported in percent and t-statistic is reported in parenthesis. To construct the vector of terrorist events we used GTD database and applied the following filters: (A) occurred outside Iraq, Syria, Afghanistan, Colombia, Somalia and Nigeria; (B) with multiple bombings or shootings and, (C) that led to more than 20 dead and/or injured. We also excluded attacks in which the event date falls into the estimated window ($t - 35, t - 16$). A total of 44 terrorist attacks resulted in the testable sample.

Panel A - CAV Responses (Means)

Period	All users
Days before:	46.56 (4.52)
Days after:	-42.03 (-6.96)
Vicinity (Whole period):	22.42 (3.27)

Panel B - CAV Responses (Evolution)

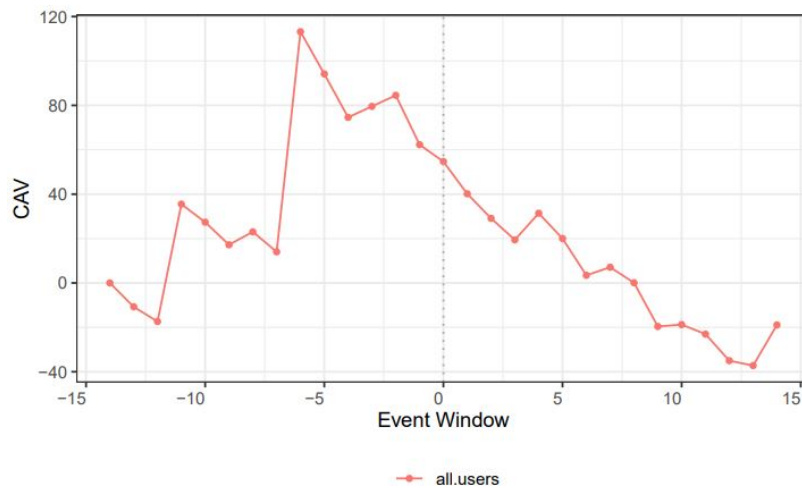


Table 3. Event study: by size This table reports the cumulative mean-adjusted abnormal volume (CAV) for the periods before and after (Panel A), and the days to the event (Panel B), across all Bitcoin services (users). The period before and after is determined in the interval [-15,-1] and [1, 15], respectively. T-Statistics is reported in between brackets. Volume is calculated as the logarithmic change of the sum of inbound and outbound volumes. Abnormal volume is mean adjusted for estimates generated in the period of 20 days before the first day in the event window. The mean-adjusted cumulative abnormal volume is reported in percent and t-statistic is reported in parenthesis. To construct the vector of terrorist events we filtered the attacks the total amount of dead and injured. The Boot-Mid test includes randomly selected (bootstrapped) attacks in which the sum of dead and injured is between 5 and 20. The Boot-Low test includes randomly selected (bootstrapped) attacks in which the sum of dead and injured is less than 5. Column Mid-Low reports the differences in the coefficients between Boot-Mid and Boot-Low. Column High-Mid reports the differences in the coefficients between High and Boot-Mid, where High coefficients are obtained from the previous results on [Table 2](#) . We also excluded events that fall into the estimated normal window (t - 35, t - 16) in each of the 50 bootstrapped runs.

CAV Responses by Size				
Period	Boot-Mid	Boot-Low	Mid-Low	High-Mid
Days before:	2.36	0.27	2.09	44.20
	(2.22)	(3.10)	(2.00)	(4.42)
Days after:	-2.47	-0.05	-2.42	-39.56
	(-1.69)	(-1.65)	(-1.64)	(-6.68)
Vicinity (whole period):	6.29	0.52	5.77	16.12
	(5.44)	(3.34)	(4.87)	(3.21)

Table 4. Event study: business groups This table reports the cumulative mean-adjusted abnormal volume (CAV) for the periods before, after and whole window (Panel A), and the days to the event (Panel B), across business groups. The period before, after and whole window is determined in the interval before [-15,-1], [1, 15] and [-15,15] days. T-Statistics is reported in between brackets. Volume is calculated as the logarithmic change of the sum of inbound and outbound volumes. Abnormal volume is mean adjusted for estimates generated in the period of 20 days before the first day in the event window. Cumulative abnormal volume is averaged over all users in each bucket of their respective portfolios. The mean-adjusted cumulative abnormal volume is reported in percent and t-statistic is reported in parenthesis.

Panel A - CAV Responses (Means)

Period	Exchange	Mixer	Gambling	Service	Dark markets
Days before:	14.96 (3.52)	24.77 (6.69)	0.32 (1.58)	-13.34 (-18.17)	-8.72 (-9.37)
Days after:	-15.30 (-7.11)	-22.78 (-11.45)	17.60 (7.06)	0.33 (4.84)	0.97 (8.61)
Vicinity (whole period):	28.70 (7.33)	24.97 (8.63)	8.71 (4.03)	-14.78 (-34.57)	-11.64 (-15.44)

Panel B - CAV Responses (Evolution)

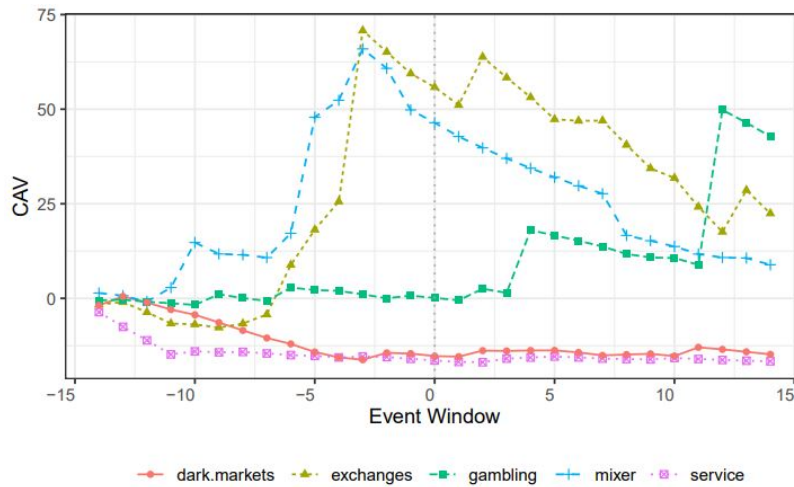


Table 5. Event study: terrorist groups This table reports the mean-adjusted cumulative abnormal volume (CAV) for the periods before, after and whole window (Panel A), and the days to the event (Panel B), across terrorists groups. The period before, after and whole window is determined in the interval before [-15,-1], [1, 15] and [-15,15] days. T-Statistics is reported in between brackets. Volume is calculated as the logarithmic change of the sum of inbound and outbound volumes. Abnormal volume is mean adjusted for estimates generated in the period of 20 days before the first day in the event window. Cumulative abnormal volume is averaged over all users in each bucket of their respective portfolios. The mean-adjusted cumulative abnormal volume is reported in percent and t-statistic is reported in parenthesis. The superscript [1] indicates placebo tests.

Panel A - CAV Responses (Means)

Period	al Qaeda	Boko Haram ^[1]	ELN ^[1]	Islamic State
15 days before:	5.72 (3.86)	-3.68 (-11.91)	-0.20 (-3.85)	5.19 (2.55)
15 days after:	0.48 (0.30)	4.08 (7.96)	2.26 (13.67)	11.86 (6.24)
the whole period:	15.35 (7.28)	-0.38 (-0.60)	1.17 (4.93)	21.07 (6.88)

Panel B - CAV Responses (Evolution)

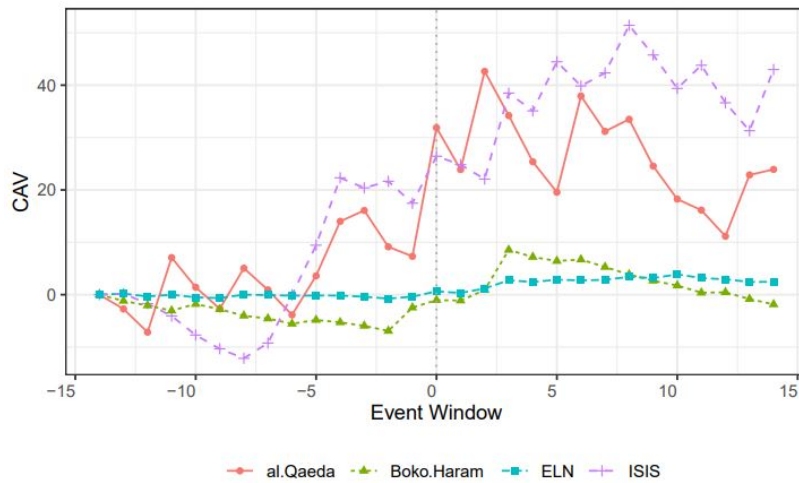


Table 6. Anomalous transfers: The table reports three-sigma rule anomaly detection results for 48 flagged users whose anomalies fall on April 2019. Columns 2 through 5 express off-sigma variables in percentage. 10% 5% and 2% of the values are off one, two and three sigmas (standard deviations) respectively. Columns 6 and 7 express the number of anomalous transfers per group of users and average per user, respectively. The last two columns express statistics in Bitcoin units.

Groups	users	$>3\sigma$	#	$\frac{\#}{user}$	mean	max
Dark markets	3.00	0.02	11.00	3.67	57.86	255.48
Exchange	24.00	0.02	65.00	2.70	2,163.24	54,320.28
Gambling	8.00	0.02	22.00	2.75	18.15	194.00
Mixer	3.00	0.01	7.00	2.33	7,581.20	26,122.56
Service	9.00	0.02	28.00	3.11	188.84	1,382.84
Total	48.00	0.02	135.00	2.81	1,481.50	54,320.28

Table 7. Ripple: The table reports statistics for transfers (A) and volume (B). Panel A reports inflows and outflows at the Gateway Ripple wallet for the first half of 2019. Flows are aggregated daily and measured in thousands of XRP (Ripple currency). Panel B reports volume frequencies for the first half of 2019 at flagged user’s Ripple wallet. Frequencies are based on five bins of XRP values. For instance, transfers less than 1,000 XRP are the most frequent occurring over 90% of the period.

Panel A: Ripple transfers		
	Inflows	Outflows
Min	34.44	31.93
Q1	64.20	58.60
Median	79.63	83.03
Mean	133.34	153.31
Q3	112.41	168.98
Max	811.50	696.35

Panel B: Volume frequency		
Bins	Obs	Frequency
Below 1,000	43,805	90.589%
1,001 to 10,000	4,051	8.377%
10,001 to 100,000	481	0.995%
100,001 to 500,000	16	0.033%
Above 500,000	3	0.006%

Table 8. Training: This table reports results for the trained set across three machine learning algorithms: Supported Vector Machine, Neural Networks and Random Forest. For each model we run the following classification formula: $Terror_{(t)} = Volume_{(t-1)} + In_{(t-1)} + Out_{(t-1)} + Life_{(t-1)} + Anonymous_{(t-1)} + Exchange_{(t-1)} + Mixer_{(t-1)} + DarkMarkets_{(t-1)} + Balance_{(t-1)} + Sigma_{(t-1)}$. Where *Terror* is an indicator variable that captures whether a terror attack occurred in that day. The vector of terrorist attacks is filtered for attacks with more than 20 injured and dead, claimed by either ISIS or al Qaeda, and that occurred in countries far away from these groups' territories. A total of 21 (9) events are allocated to the training (validation) set. *Volume* is the logarithm of the sum of inbound and outbound transfers measured in Bitcoin units. *In* is the logarithm of the inbound transfer value measured in Bitcoin units. *Out* is the logarithm of the outbound transfer value measured in Bitcoin units. *Life* is the life of the wallet, defined as the distance in days from the first trade and the last trade, associated with sending or receiving the funds from users in our sample. *Anonymous* is an indicator variable that captures whether funds came from, or were sent to, an anonymous address. *Exchange* is an indicator variable that captures whether funds came from, or were sent to, an exchange address. *Mixer* is an indicator variable that captures whether funds came from, or were sent to, a mixer address. *DarkMarkets* is an indicator variable that captures whether funds came from, or were sent to, a dark markets address. *Balance* captures the total balance of the user's wallet in Bitcoin units. *Sigma* is an indicator variable that captures whether transfer size falls off three standard deviations from user's last three months historical mean. All predictors used in the classification model are at the transfer-user level and lagged in one day.

	Supported Vector Machine	Neural Networks	Random Forest
Accuracy	0.9939	0.9939	1
Kappa	0.1076	0.0895	1
No Information Rate	0.9937	0.9937	0.9937
P-Value (ACC > NIR)	0.0525	0.0949	<0.0001
McNemar's p-value	<0.0001	<0.0001	<0.0001
Prediction (NO/NO)	285210	285216	285247
Prediction (NO/YES)	37	31	0
Prediction (YES/YES)	106	87	1812
Prediction (YES/NO)	1706	1725	0
Balances Accuracy	0.5292	0.5240	1
Precision	0.9941	0.9940	1
Recall	0.9999	0.9989	1
F-measure	0.9969	0.9969	1

Table 9. Validating: The table reports confusion matrix and statistics for the validation set. The first columns report statistics for the flagged user. The remaining columns report mean statistics for group of users. The validation set is determined on anomalous transfers that occurred after January 1, 2018. Results in this table are for a 10-fold cross-validated random forest model with 500 trees and nine variables at each split.

	User	Exchange	Gambling	Service
Accuracy	0.627	0.752	0.815	0.701
95% Confidence Interval	(0.626, 0.629)	(0.742, 0.756)	(0.812, 0.818)	(0.695, 0.705)
Sensitivity	0.628	0.755	0.819	0.704
Specificity	0.424	0.246	0.25	0.277
Prevalence	0.996	0.996	0.995	0.995
Detection Rate	0.626	0.751	0.815	0.700
Detection Prevalence	0.628	0.755	0.818	0.704
Balanced Accuracy	0.526	0.500	0.534	0.490
Users	1	93	49	55

Appendix A1. Lexicon: This table describes the main terms used in the paper by order of appearance.

General

BTC	the cryptocurrency Bitcoin.
Dark markets	a portfolio of dark market users. Users in this group provide services of usually selling illicit products or services.
Exchange	a portfolio of exchange users. Users in this group provide exchange services such as conversion fiat to crypto, crypto to fiat and crypto to crypto currencies.
Gambling	a portfolio of gambling users. Users in this group provide gambling services such as casino and bet online platforms.
ICO	an Initial Coin Offering (ICO) is a method of raising funds through the use of cryptocurrencies. Its use is most popular in projects that have not yet fully developed their blockchain platform, product, or service. The payment is usually made with Bitcoin or Ethereum, but in some cases, fiat currency is also accepted.
Mining	a portfolio of mining users. Users in this group mine Bitcoin and/or other cryptocurrencies.
Mixer	a portfolio of mixer users. Users in this group provide tumbling services. The service consists of reshuffling cryptocurrencies into hundreds of transactions and interpolating transactions with other users to decrease or eliminate traceability.
Service	a portfolio of service users. Users in this groups provide general Bitcoin services such as online payments, transfers and cold storage.
Wallet	the wallet holds user’s Bitcoin balance but does not actually contain the bitcoins. It stores the user’s Bitcoin address and private key to access the Bitcoin blockchain. When users make payments, the wallets use the key to digitally sign the transactions proving ownership of their coins in the network.
XRP	the cryptocurrency Ripple.

Anomalous Transfers

$>\sigma$	transfers larger than one standard deviation.
$>2\sigma$	transfers larger than two standard deviation.
$>3\sigma$	transfers larger than three standard deviation.
#	number of anomalous transfers.
#/users	number of anomalous transfers per group of users.

Random Forest

Suppose a 2x2 table with notation:

	Realized		
Predicted	Event	A	B
No Event	No Event	C	D

A	the model predicts the event occurring and the event occurs (correct prediction).
B	the model predicts the event occurring but the event does not occur (false positive outcome).
C	the model predicts the event not occurring but the event occurs (false negative outcome).
D	the model predicts the event not occurring and the event not occur (correct prediction).
Balanced Accuracy	the equally weighted accuracy between sensitivity and specificity or (sensitivity + specificity)/2.
Detection Prevalence	the proportion of detected events out of total outcomes $(A+B)/(A+B+C+D)$.
Detection Rate	the proportion of detected true events out of total outcomes or $A/(A+B+C+D)$.
OOB estimate rate	the OOB (out-of-bag) estimate of error rate is a useful measure to discriminate between different random forest classifiers.
Kappa	the kappa statistic is a metric that compares an observed accuracy with an expected accuracy (random chance).
Sensitivity	the correct proportion of events occurring $A/(A+C)$.
Specificity	the correct proportion of events not occurring $D/(B+D)$.