

Coins for Bombs

The Predictive Ability of On-chain Transfers for Terrorist Attacks

Amiram, Dan; Jørgensen, Bjørn N.; Rabetti, Daniel

Document Version

Final published version

Published in:

Journal of Accounting Research

DOI:

[10.1111/1475-679X.12430](https://doi.org/10.1111/1475-679X.12430)

Publication date:

2022

License

CC BY-NC-ND

Citation for published version (APA):

Amiram, D., Jørgensen, B. N., & Rabetti, D. (2022). Coins for Bombs: The Predictive Ability of On-chain Transfers for Terrorist Attacks. *Journal of Accounting Research*, 60(2), 427-466. <https://doi.org/10.1111/1475-679X.12430>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 04. Jul. 2025

Coins for Bombs: The Predictive Ability of On-Chain Transfers for Terrorist Attacks

DAN AMIRAM , * BJØRN N. JØRGENSEN , †
AND DANIEL RABETTI 

Received 30 November 2020; accepted 3 March 2022

ABSTRACT

This study examines whether we can learn from the behavior of blockchain-based transfers to predict the financing of terrorist attacks. We exploit blockchain transaction transparency to map millions of transfers for

*Coller School of Management, Tel Aviv University; †Copenhagen Business School and Hanken School of Economics

Accepted by Luzi Hail. We are grateful to an anonymous reviewer, Sanjay Banerjere, John Barrios, Thomas Bourveau, Marie Briere (discussant), Brian Burnett, Hans Christensen, Atif Ellahie, Vivian Fang, Sivan Frenkel, Jeffrey Hoopes, Eva Labro, Christian Leuz, Shai Levi, Tsafir Livne, Evgeny Lyandres, Daniele Macciocchi, Mark Maffett, Maximilian Muhn, Jacob Oded, Kasper Regenburg, Steve Rock, Sugata Roychowdhury, David Schoenherr (discussant), Daniel Scott Cohen, Harald Uhlig, Tsahi Versano, Regina Wittenberg Moerman, Avi Wohl, Andrew Wu (discussant), Anastasia Zakolyukina, Yanlei Zhang, Peter Zimmerman (discussant), and participants at the 2021 *Journal of Accounting Research* Conference, International Association for Quantitative Finance (IAQF), 7th Fin-Fire Conference on Challenges to Financial Stability, Crypto and Blockchain Economics Research Conference, 3rd Bergen FinTech Conference, 4th Shanghai-Edinburgh Fintech Conference, 6th Fintech International Conference, U.S. Department of the Treasury, Copenhagen Business School, and Tel Aviv University, for invaluable comments. We thank WhiteStream for providing intel on the Sri Lanka Easter bombing and DeepSeek for additional insights into terrorist operations on the dark web. The authors are thankful to the Henry Crown Institute, the Danish Finance Institute, and the Coller Blockchain Research Institute for financial support. An online appendix to this paper can be downloaded at <http://research.chicagobooth.edu/arc/journal-of-accounting-research/online-supplements>.

March 2022

hundreds of large on-chain service providers. The mapped data set permits us to empirically conduct several analyses. First, we analyze abnormal transfer volume in the vicinity of large-scale highly visible terrorist attacks. We document evidence consistent with heightened activity in coin wallets belonging to unregulated exchanges and mixer services—central to laundering funds between terrorist groups and operatives on the ground. Next, we use forensic accounting techniques to follow the trails of funds associated with the Sri Lanka Easter bombing. Insights from this event corroborate our findings and aid in our construction of a blockchain-based predictive model. Finally, using machine-learning algorithms, we demonstrate that fund trails have predictive power in out-of-sample analysis. Our study is informative to researchers, regulators, and market players in providing methods for detecting the flow of terrorist funds on blockchain-based systems using accounting knowledge and techniques.

JEL codes: G15, G18, G29, K29, K42, M40, M41, O16

Keywords: transparency; terrorist financing; economics of blockchain; forensic accounting; bitcoin

Right now, large parts of the field of crypto are sitting astride of—not operating within—regulatory frameworks that protect investors and consumers, guard against illicit activity, ensure for financial stability, and yes, protect national security.

—Gary Gensler (August 2021), chair of the U.S. Securities and Exchange Commission

1. *Introduction*

In recent years, international money transfers through blockchain-based currencies have grown significantly.¹ The proliferation of cryptocurrencies has had two opposing effects related to the transparency of the international money-transfer system. On the one hand, governments, financial institutions, and market regulators invest hundreds of millions of dollars and numerous person hours in curtailing illegal transfers through the traditional financial system (Belasco et al. [2018]). Cryptocurrencies and, in particular, Bitcoin, the most popular 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 technical knowledge of the blockchain system and the ability to analyze information in public blockchain ledgers (ICAEW [2018]). The increased transparency of transfers may help outsiders identify and predict illicit activities by monitoring abnormal transactions.

¹ As of February 2022, more than 17,000 crypto assets (e.g., blockchain-based currencies) are listed in more than 450 crypto exchanges across the globe, totaling 1.7 trillion dollars in market capitalization (<https://coinmarketcap.com/>).

This study examines whether we can learn from the behavior of blockchain-based transfers in the vicinity of large-scale, highly visible terrorist attacks to predict similar events.² Many Bitcoin transfers have been attributed to illegal activities (Foley, Karlsen, and Putnins [2019]), and cryptocurrencies can be used by terrorists for donation campaigns (Irwin and Milad [2016], Dion-Schwarz, Manheim, and Johnston [2019]). As the analysis of donation campaigns (i.e., transfers between donors and terrorist organizations) depends on leaked addresses, our focus is on terrorist attacks (i.e., transfers between the terrorist organization and operatives on the ground), which do not require foreknowledge of terrorist wallets. The transparency of the blockchain system suggests that it may be possible to identify transfers associated with terrorist attack financing. Financing attacks require intense money laundering, entailing abnormal volume in the vicinity of an event as a potential proxy for terrorist financing.³ Our methods may help outsiders detect funds associated with terrorist attacks and, perhaps more importantly, help them predict these events.⁴

Additional considerations motivate this study. First, the rise of blockchain analytics provides several avenues for academic research. Fund transfers on the Bitcoin blockchain are in the public domain and can be traced back to wallets.⁵ With the development of deanonymizing algorithms, such as the one we employ here, it is possible to connect additional addresses by checking whether a known address co-spent with other addresses on the chain.⁶ At the end of the mapping, we account for thousands of addresses at the users' wallet level, their respective balances at a given point in time, and the flow of funds throughout time. As a result, blockchain-mapped data provide a unique setting for the analysis of transfers and participants' interactions.⁷

² Our study focuses on Bitcoin, as it is the most liquid and used public blockchain-based cryptocurrency. Bitcoin recently surpassed 1 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 a detailed description of the Bitcoin protocol).

³ We use the term "vicinity" as a measure of chronological closeness (i.e., the days surrounding the date of the event).

⁴ This method may be applied to other illicit activities, such as laundering funds for tax evasion, financing political unrest, and large-scale purchases of goods (e.g., weapons) on black markets.

⁵ Wallet owners are anonymous (although their activity/fund transfers are public), except for some wallet owners who reveal their identity willingly or by accident. For instance, we focus on users that are business entities in the chain. These users willingly provide wallet addresses to their clients.

⁶ By co-spending, we mean that two or more addresses share the same input in one transaction, which suggests that they belong to the same user. See subsection 3.1 for more details. Moreover, the computer science literature develops refined techniques that allow parties to identify users' IP addresses in some cases (e.g., Kang et al. [2020]). Because users' identities may be revealed through these techniques, Bitcoin is considered pseudo-anonymous.

⁷ Unlike in public blockchain systems, in traditional banking fund transfers are opaque to outsiders, including to a large extent, regulators. Additionally, even regulators generally

Additionally, individuals with blockchain knowledge and (arguably) accounting expertise may have a relative advantage in analyzing blockchain transactions. The Institute of Chartered Accountants in England and Wales (ICAEW), one of the oldest accounting organizations globally, has developed a conceptual view in which accounting expertise (broadly speaking) is helpful to assess blockchain technology, as the blockchain can be viewed as an accounting system that maintains a ledger of accurate financial information and provides clarity over ownership of assets.⁸ In this study, we use accounting knowledge to investigate the financing of terrorist attacks in blockchain-based systems.

Large-scale on-the-ground terrorist attacks require financing to buy weapons, explosives, and other equipment and to pay operatives and their families. If financiers are worried about eliminating their traces in the blockchain, they will likely try to launder their transactions. This consists 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 terrorist attack, say, to buy machine guns on the black market. If this amount is reshuffled 100 times, over a million dollars is generated in volume due to laundering. We thus argue that one would observe abnormally large volumes of transfers in the vicinity of terrorist attacks.⁹ Using abnormal volume to investigate an event's information content has been common since the work of Beaver [1968].

To conduct the empirical analysis, we aggregate millions of transactions for hundreds of users in the Bitcoin blockchain and classify them into six groups: dark markets, exchanges, gambling platforms, mining, mixers, and other services. Users of dark markets usually sell illicit products or services on the dark web. Exchanges provide services such as converting fiat-to-crypto, crypto-to-fiat, and crypto-to-crypto. Gambling platforms operate online casinos and betting markets. Miners mine Bitcoin and other cryptocurrencies. Mixers provide tumbling services, which consist of reshuffling cryptocurrencies into hundreds of transactions and interpolating transactions with other users to decrease or eliminate traceability. The services group provides general Bitcoin services, such as online payments, transfers, and cold storage. We merge these data with terrorist attack data compiled

depend on the information they receive from banks by law or court rulings to gain insights into transfers. Beyond that, the global regulatory system is fragmented, and regulators can generally follow only the parts of the fund trail that are under their jurisdiction.

⁸ The ICAEW describes blockchain as “an accounting technology. It is concerned with transferring 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.... For accountants, using blockchain provides clarity on ownership of assets” (ICAEW [2018]).

⁹ Furthermore, as we elaborate below, even if the documented abnormal volume is not necessarily driven by terrorist wallets but by other market players screening for abnormal activity, their association with the event is in itself informative.

from the comprehensive list of attacks from the Global Terrorism Database (GTD) for 2015 to mid-2019.

The vast majority of GTD events are not expected to be relevant for our study because they consist of small-scale attacks (i.e., assassinations, hostage taking, infrastructure attacks, unarmed assaults, arson, stabbings, and melee attacks) concentrated mostly in regions marked by local insurgencies, such as Afghanistan, Iraq, and Syria. However, to provide a benchmark for our examinations, we keep these small events in our baseline specification, resulting in 21,323 unique events.¹⁰ For the main specification, our focus is on large-scale, highly visible events, as they are more likely to require financing.

We identify large-scale events based on whether a terrorist attack has been highly publicized in the media. To do so, we match the GTD data set with Wikipedia articles.¹¹ Wikipedia uses a series of criteria to determine event notability (i.e., whether it deserves a standalone Wikipedia article). For instance, a standalone article should be based on several reliable published sources.¹² In line with our focus on large-scale events, the average number of victims is approximately five times larger in a notable attack (26) than in the baseline specification (5). Finally, as notable events occur mainly within terrorists' home territories and Bitcoin is more likely to finance extra-territorial activities (e.g., Irwin and Milad [2016] and Dion-Schwarz et al. [2019]), we further filter for attacks in foreign territories in our third specification. This last specification is used only in the predictive section of the paper to improve the model's accuracy.

We start our analysis by examining Bitcoin's cumulative abnormal volume (CAV) responses in the vicinity of terrorist attacks. Notable attacks often entail multiple bombings (77.47%) or mass shootings (21.54%) and involve many operatives on the ground; therefore, these events are likely to require substantial financing. If terrorists use Bitcoin to finance these events, they likely use mixers or mixer-like transactions to camouflage their transfers. If executed in the vicinity of the event, such activities could be observed via abnormal volume analysis. Consistent with this prediction, we find a sharp increase in CAV in the period preceding notable attacks (9.53%), followed by a decrease in the weeks thereafter (7.33%). This result is robust to controlling for events likely to affect Bitcoin on-chain volumes, such as hardforks, price peaks, and public holidays.¹³ We obtain the

¹⁰ GTD fragments events. For instance, the Sri Lanka Easter bombing is treated as eight events, because of multiple bombings in different locations. We treat all these attacks as one unique event.

¹¹ This approach has been shown to produce more informative events (e.g., Xia and Gu [2019]).

¹² The Sri Lanka Easter bombing standalone Wikipedia article is based on over 300 sources, including *The New York Times*, *The Independent*, *Reuters*, *The Washington Post*, and *The Guardian* (https://en.wikipedia.org/wiki/2019_Sri_Lanka_Easter_bombings).

¹³ "Hardfork" (or hard fork) is a blockchain split that leads to a new blockchain starting from the split block. For instance, Bitcoin Cash is a Bitcoin blockchain hardfork created on

reported CAV responses by bootstrapping notable events and reporting the mean coefficients to mitigate selection concerns. In addition to the main specification of notable events, we also examine the baseline specification containing thousands of small-scale events. This procedure not only helps us address our second prediction that Bitcoin is more likely to finance large-scale events, but also works as a placebo test to mitigate model specification concerns. We find that the bootstrapped mean-CAV response for the baseline specification is small and statistically insignificant in the weeks before a small-scale event occurs.

We then analyze CAV responses across different users by splitting users into groups of similar services. We expect abnormal volume, in the vicinity of the event, to concentrate at crypto exchanges and mixers, as operatives on the ground need these services to withdraw Bitcoin, make online payments and camouflage illicit activity. We find that users of exchanges and mixers present large positive CAVs in the weeks preceding the events. A significant variation in the quality of regulation of crypto exchanges and the level of compliance procedures, such as know-your-customer (KYC) and anti-money laundering (AML), potentially facilitates illegal activities (Amiram, Lyandres, and Rabetti [2021]).¹⁴ As the flow of illicit funds is more likely to go through unregulated exchanges than through regulated ones, we follow Cong et al. [2021] and separate these two groups according to whether they are licensed by U.S., U.K., or Japanese authorities. We find that unregulated exchanges drive most of the CAV responses. Finally, due to sudden demand for mixing services, other services that commonly rely on mixers for their daily activities are expected to have lower than expected volume because these services are interconnected on the blockchain. Consistent with this prediction, the abnormal volume in the vicinity of notable attacks is smaller than expected for dark markets, gambling, and other services.

Collectively, the event-study results are consistent with terrorists using money laundering techniques to finance attacks primarily through unregulated exchanges and mixers. However, our research design precludes us from establishing a direct relation to terrorist organizations. The observed abnormal volume may reflect transfers from nonterrorist entities with knowledge of the mechanics of money laundering or some other information on pending attacks. For instance, the laundering process associated with the attack may result not from direct financing but from terrorist groups anticipating increased scrutiny of the blockchain by authorities after a large-scale attack. This caveat does not alter the inferences we draw on the

August 1, 2017, to increase Bitcoin capacity and block size. See the online appendix for details on these events.

¹⁴ KYC is the process of verifying a customer's identity. AML is a series of procedures to identify and prevent money laundering. Together, they provide the highest level of compliance in crypto exchanges to prevent illegal activities, such as money laundering, terrorist financing, and tax evasion.

association between abnormal volume and events and on the informative role of abnormal volume in predicting large-scale terrorist attacks.

We complement our main analysis with a case study-type examination of the Sri Lanka Easter bombing.¹⁵ In addition to being the largest attack in a foreign territory for which the Islamic State of Iraq and Syria (ISIS) claimed responsibility, it involved extensive logistics (i.e., multiple bombings and many terrorists on the ground), and several news sources claimed that Bitcoin was used to finance it.¹⁶ We identify suspicious users with abnormal transfers in the vicinity of the event by applying forensic accounting and anomaly detection techniques. We then examine flagged users for whether alternative channels explain abnormal transfers around the attack and end up with one user for which no alternative explanations are found. Focusing on this suspicious user, we assess whether addresses linked to its wallet have a history of association with other crimes. We track transfers backward in time and find evidence of associations with several reported crimes, including funding of jihadi cells in Syria. We also track the funds transferred after the Sri Lanka attack and find that some of these funds were likely converted 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 in reserves.¹⁷ These findings provide further insights into the mechanics of terrorist crypto financing.

Finally, we implement machine-learning algorithms to predict terrorist attacks. We consider three algorithms, the supported vector machine (SVM), neural networks (NNs), and random forest (RF), as they stand out as state-of-the-art solutions for supervised nonlinear learning classifiers. We train these models across all users in the data set; however, the user flagged with anomalous transfers around the Sri Lanka bombing provides the best training accuracy. The list of notable terrorist events is the same as in the main specification but adjusted for foreign attacks and filtered for Bitcoin price peaks and hardforks, resulting in 30 events mostly claimed by either ISIS or al Qaeda. We split the data into training and validation sets, the former (latter) comprising approximately 70% (30%) of the observations and 21 (9) events. Although all three models have predictive power, the random forest model achieves the best performance for predicting out-of-sample terrorist attacks a day before they occur. These results are potentially insightful to market participants and agencies concerned with the on-chain flow of terrorist funds. However, a couple of caveats apply, including one

¹⁵ According to *The New York Times*, the attack killed over 250 people (45 children) and wounded over 500 people (<https://www.nytimes.com/2019/04/24/world/asia/sri-lanka-easter-bombing-attacks.html>).

¹⁶ See, for instance, <https://en.globes.co.il/en/article-exclusive-isis-funded-sri-lanka-bombings-with-bitcoin-donations-1001284276>.

¹⁷ Although we followed the trails of funds likely associated with the Sri Lanka Easter bombing, proving that this big wallet belongs to ISIS is beyond the scope of this paper.

related to the limited information we have on terrorist wallets. Additionally, our results are limited to a group of large Bitcoin service providers (users). Finally, our results may apply only to Bitcoin-like ecosystems.

This paper provides several contributions to the literature. First, it joins studies examining the predictive power of accounting information and knowledge in nonaccounting settings. For instance, Godsell, Welker, and Zhang [2017] test whether accounting models can detect earnings management in an import relief setting. The strategy of using accounting knowledge to obtain information in other settings is not limited to contemporaneous work. Accounting information has already been used in classic studies to predict stock returns (Bernard and Thomas [1989], Sloan [1996], Abarbanell and Bushee [1997], Piotroski [2000]), bankruptcy (Altman [1968], Beaver [1968], Ohlson [1980], Zmijewski [1984]), and fraud (Singleton and Singleton [2010], Kranacher and Riley [2019]). The uniqueness of our study lies not only in the novelty of the source of “accounting information,” but also in the unobservability of the ultimate owner; therefore, predicting the ultimate owner (i.e., a terrorist organization) is also part of our research goal and contribution.

Second, our study contributes to the relatively new literature examining transparency in alternative financing platforms. For instance, Michels [2012] examines how voluntary disclosures attenuate market inefficiencies in peer-to-peer lending markets. Bourveau et al. [2021] examine crypto analysts’ role in assessing the quality of initial coin offerings. While the interplay of disclosure and consumer regulation has been studied by Cascino, Correia, and Tamayo [2019] in reward crowdfunding markets, Lyandres, Palazzo, and Rabetti [2021] examine the effects of disclosed information on crypto-based projects’ operational and financial performance. We show how transaction transparency enables outsiders to study the flow of funds and users’ interactions in a blockchain-based financing system.

Finally, this paper contributes to the literature that examines the methods of terrorist organizations (e.g., Pieth [2002], Schott [2006], and Rudner [2010]). Specifically, our paper provides novel empirical evidence that responds to studies questioning whether cryptocurrencies contribute to terrorist financing (e.g., Irwin and Milad [2016] and Dion-Schwarz et al. [2019]). In this regard, our findings could also inform the current debate over cryptocurrency regulation (Foley et al. [2019], Fusaro and Hougan [2019], Griffin and Shams [2020], Makarov and Schoar [2020], Sokolov [2021], Amiram et al. [2021], and Cong et al. [2022]) by demonstrating that large unregulated market institutions might unwillingly operate as facilitators of terrorism.

2. *The Financing of Terrorism*

The financing of terrorism can be costly because it involves not only executing an attack, but also establishing, growing, and maintaining operations in distant territories (Malkin and Elizur [2002]). Research has shown

that following money trails enables the tracking of terrorist attacks. For instance, tracking terrorist funds, Limodio [2021] shows that terrorist organizations become more active after a positive funding shock. Another example is a widely used informal system, Hawala, which usually involves a sender, a receiver, and two dealers.¹⁸ As this system does not involve the transfer of funds across countries, it is very difficult to track (Buencamino and Gorbunov [2002]). In addition, the use of a variety of codes provides security for the transaction (Kiser [2005]). Al Qaeda, for instance, for years circulated funds under the authorities' radar through Hawala before the September 11, 2001, attack. In addition to cash transfers, terrorists also accumulate gold reserves. Al Qaeda and the Taliban used a Halawa network to move millions of dollars worth of gold around the world (DeYoung and Farah [2002], and Ashcroft and Snow [2003]).

However, following the September 11 attacks, authorities increased scrutiny and shut down several Hawala operators worldwide (Heng and McDonagh [2009] and Navias [2002]).¹⁹ Additionally, because not all Hawala dealers operate illicitly, an internal movement attempted to regulate the system (Borgers [2009]). As a result, terrorists started seeking options beyond Hawala. These include more formal remittance systems, such as the Western Union, eBay, and PayPal (Cook and Smith [2011]). Although the evidence suggests that terrorists jointly use Hawala, as well as formal remittance systems, the authorities' main focus remains on informal networks (Acharya [2012]).

Cryptocurrencies are notorious for their potential criminal use due to the pseudo-anonymity of their ownership, widespread usage, and the lack of regulation of crypto exchanges and other services. Bitcoin, the most popular cryptocurrency, is heavily used in dark markets, with an estimated one-half of its transactions and one-fourth of its users linked to illicit activities (Foley et al. [2019]). Silk Road, the most prominent dark market, at its peak of activity, was moving hundreds of millions of dollars worth of cryptocurrencies and providing services to over 100,000 buyers until the Web site was shut down by the US Federal Bureau of Investigation and the owner arrested.²⁰ Several other marketplaces similar to the Silk Road were launched and eventually seized; however, some dark markets are still operating.

Cryptocurrencies are also used as the preferred means of payment by cyber attackers. Bitcoin alone has over 224,000 addresses listed in illicit

¹⁸ The sender provides the local Hawala dealer with funds. The local Hawala dealer contacts a corresponding Hawala dealer in the country where the sender wants to send the funds. The Hawala dealer in the target country then provides the corresponding amount to the receiver.

¹⁹ See, for instance, <https://archives.fbi.gov/archives/newyork/press-releases/2010/nyfo091510a.htm> or <https://www.justice.gov/archive/usao/nys/pressreleases/January10/safarharezactalarrestspr.pdf>.

²⁰ See <https://www.justice.gov/usao-sdny/pr/senior-adviser-operator-silk-road-website-pleads-guilty-manhattan-federal-court>.

address-screening services, such as *Bitcoin Abuse*. The number of “sextortion” and ransomware attacks have increased with Bitcoin’s growing popularity.²¹ A recent case, the ransomware attack at Colonial Pipeline, spurred U.S. authorities to launch a task force to fight cyberattacks. At Colonial, the attacker was able to raise \$4.5 million in Bitcoin payments. However, due to blockchain transparency, agencies and other outsiders tracked the money and recovered some of it.²² Nonetheless, most paid cyberattackers succeed in keeping the funds.²³

Mounting evidence of Bitcoin use in financing illicit activities, such as dark markets and cyberattacks, suggests its potential use by terrorists (Irwin and Milad [2016]). Although this argument may be contested (Dion-Schwarz et al. [2019]), there is strong evidence from as early as 2013 of cryptocurrencies in terrorists’ donation campaigns.²⁴ In arguably the most prominent such campaign to date, Hamas’s military arm, the al Qassam Brigades, used its main Web page, social media channels, and officials to call for donations following an Israeli embargo on external financial aid (Katisiri [2019]). However, the US Department of Justice (DOJ) recently dismantled the Bitcoin campaigns of the al Qassam Brigades as well as those of al Qaeda and ISIS. This has hindered terrorist groups from raising funds via cryptocurrencies. Chainalysis, a blockchain analytics provider, reports that donation campaigns were able to raise only negligible amounts and are just a tiny fraction of the total flow of illicit funds.²⁵

Given the extensive evidence of the use of cryptocurrencies in illicit activities, including financing terrorist donation campaigns, terrorists may also use them to finance their attacks. However, as terrorist wallets involved in attack financing are not revealed to the public, we use event-study and accounting techniques to descriptively examine Bitcoin transactions on services likely used by terrorists in the financing process.

²¹ “Sextortion” is a cyberattack where the victim is asked to pay cryptocurrencies to prevent the attacker from disclosing information related to the victim’s usage of adult web services. Ransomware is an attack that usually targets corporations where hackers have stolen sensitive data, such as clients’ private information, and demand payments in exchange for not leaking the data to the public.

²² Agencies could follow the money trail on the blockchain and recover a large part of it (<https://www.cnbc.com/2021/06/07/us-recovers-some-of-the-money-paid-in-the-colonial-pipelineransom-officials-say.html>).

²³ Cryptocurrency-based crime hit a new all-time high in 2021, with illicit addresses receiving \$14 billion over the course of the year (<https://blog.chainalysis.com/reports/2022-crypto-crime-reportintroduction/>).

²⁴ See the online appendix for a detailed discussion of donation campaigns.

²⁵ See report in <https://blog.chainalysis.com/reports/cryptocurrency-crime-2020-report>.

3. Data

3.1 BLOCKCHAIN DATA ACQUISITION AND CLASSIFICATION

The pseudo-anonymity of Bitcoin imposes nontrivial challenges to the identification of market participants. Owners of user wallets must be identified off-blockchain, and each wallet may contain hundreds, even thousands, of addresses. 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 additional information is revealed. The connection between wallet ownership and the identity of a real person or business entity may become publicly known because owners registered on an exchange were leaked to social media or other Web sites 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 obtains access to several addresses from Bitcoin providers (e.g., exchanges), enabling those addresses to be revealed. Second, once the addresses are obtained, data fusion algorithms are employed to associate several connected addresses with a specific user.

To aggregate addresses to the user level, we implemented the union-find algorithm, which has been used in several academic applications (e.g., Kappos et al. [2018], Tasca, Hayes, and Liu [2018], and Foley et al. [2019]).²⁶ 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 et al. [2018], Foley et al. [2019]). Based on this algorithm, we identify the wallets of 343 Bitcoin users.²⁷

A user in our data set is a collection of wallets for a business entity and may contain thousands of transactions and addresses. For instance, approximately 200,000 addresses and over 800,000 transactions linked to the Binance exchange's wallets are treated as one user in our data set. As our data include the main service providers in the Bitcoin ecosystem, we likely cover most of the transactions by known participants in the

²⁶ The algorithm was designed by Cormen et al. [2001] and, in the context of cryptocurrencies, was 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>.

²⁷ Consider the following as a simplified explanation of how the algorithm works. Say the addresses X and Y co-spend (e.g., if the transaction is 1,000, each spends 500 units) in the transaction Alpha, and the addresses Y and Z co-spend in the transaction Beta. As Y appears in two different transactions by co-spending separately with X and Z, all addresses (X, Y, and Z) 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 the user data level.

sample period.²⁸ To classify Bitcoin users into business groups, we follow prior research.²⁹ We classify users into five portfolios: dark markets, exchange, gambling, mixer, and service, where *Dark markets* represent online shops, usually on the dark web, responsible for selling illicit goods. *Exchange* contains crypto exchanges responsible for providing exchange and withdrawal services. *Gambling* includes gambling, rewarded games, and betting platforms. *Mixer* relates to business associated with reshuffling cryptos into several transactions and wallets to eliminate traces in the blockchain. *Service* includes general business transactions, such as online payments, credit cards, and wallet storage.

3.2 BLOCKCHAIN DATA PROCESSING AND STATISTICS

The data collected present a series of challenges. First, some users become inactive. For instance, the Silk Road became inactive after the FBI seized it. This translates to user sample variation around terrorist attacks, with some users absent around the most recent attacks. Second, even among active users, some 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. We therefore 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 169,582 cold storage addresses, mostly of exchanges and services, that are irrelevant to our analysis, as they are used only to store coins. Finally, we exclude transactions that had no interaction with other users (i.e., no output). The total number of transactions excluded represents only a tiny fraction (0.55%) of the data obtained.³⁰ The final data set contains 342 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. Users have a total of 98.95 million addresses resulting from data fusion of 135.75 million transactions from 2014 to 2019. The extensive set of addresses and the total number of transactions indicate the complexity of mapping the 342 large users in this data set. *Exchange* and

²⁸ For comparison purposes, Foley et al. [2019] exclude users for which the transfer of cash does not involve the acquisition of goods or services. Although they do not mention the number of users, they state that 88.4 million transactions were excluded. We obtain 105.67 million transactions for the same group of users, which means that our sample is approximately 20% larger, even though we cover a shorter time (2014–2019) than they do (2009–2017).

²⁹ We follow the <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 markets* as a category. Classifying users into groups is also present in the accounting literature. For instance, Bushsee [1998] assigns investors to dedicated, transient, and quasi-indexing investor groups to examine their influence on firms' R&D investment.

³⁰ Our inferences are unchanged if these transactions are included in the final sample.

TABLE 1
Blockchain Sample Construction

Panel A: Sample construction (user)								
	Users	Transactions	Address	Volume	Average	Median	Balance	Life
Obtained:	343	136.5	99.79	222.24	21.90	0.01	1,429.57	1,287
Excluded:								
Mining	1	0.61	0.67	0.03	78.07	15.20	0.06	2,220
Cold Storages	–	0.13	0.17	0.01	658.03	1.74	83.86	979
Lower Interaction	–	0.01	0.01	0.01	182.03	45.11	0	125
Used:								
Dark Markets	99	11.26	7.27	31.21	23.69	0.01	10.3	1,583
Exchange	101	62.12	51.99	155.31	7.85	0.02	4,884.85	1,496
Gambling	50	18.53	6.64	7.78	0.78	0.01	2.83	1,345
Mixer	36	0.34	0.28	0.37	6.13	0.01	1.55	81
Services	56	43.5	32.48	23.12	2.00	0.01	165.61	1,287
Total	342	135.75	98.95	217.8	10.24	0.01	1,458.19	1,290
Panel B: Aggregate summary (user)								
	Mean (\$)			SD (\$)		Median (\$)		
Volume (billions)	1.52			7.07		0.01		
Balance (millions)	4.35			20.10		0.03		
Fee (thousands)	0.70			4.30		0.01		
Panel C: Distribution of volume (daily)								
Mean (K\$)	SD (K\$)	Median (K\$)		Skewness (#)		Kurtosis (#)		
73.55	61.07	58.45		2.03		6.70		
Q1	Q2	Q3		Q4		Q5		
0.00	30.41	58.45		98.85		544.17		

This table reports descriptive statistics on the composition of the blockchain sample. Panel A reports the number of users, transactions, addresses, and volume for the obtained, excluded, and used blockchain samples. We report variables in units, except Transactions, Address, and Volume, which are in millions, and Life, which is in days. Panel B reports aggregated user-level statistics in dollars. Panel C reports the distribution of daily volume in dollars, including quintiles.

Service have the largest number of addresses across groups, with 51.99 million and 32.48 million, respectively. However, when volume is considered, *Dark markets* become the second largest group, with 31.21 million Bitcoins. *Exchange* has the largest number of Bitcoins transferred, with 155.31 million units, or over 70% of transfers, in volume during the period. Although some transfers contain large numbers of Bitcoins (mean of 10.24), most of the transfers are very small (median of 0.01) across all users in the data set. Notably, users in the *Dark markets* group have the largest number of Bitcoins transferred per transaction (mean of 23.69), suggesting that this type of business accumulates Bitcoins before cashing them out or converting them to other cryptocurrencies. *Exchange* users tend to maintain a larger balance of Bitcoins than users in different business niches. This 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 Markets* had the longest average life (4.3 years) among all groups. The reason is that these were the first type of business to adopt Bitcoin. In contrast, *Mixer*, a service that receives dirty coins (coins linked to illicit activities or additional privacy needs), mixes them into several transactions and aggregates them into clean coins (difficult-to-trace coins), has the shortest average life (2.7 months), consistent with this type of service reducing traceability.

Panel B reports summary statistics in dollars aggregated at the user level. The average user moves 1.52 billion dollars worth of Bitcoins during its lifetime. The standard deviation among users is 7.07 billion dollars, which indicates that a few large services dominate the market. A median of 0.01 billion dollars shows that most users move small amounts, reflecting that Bitcoin is still young (i.e., most of its services have not yet matured).

Panel C reports the distribution of daily volume in dollars. An average day moves 73.55 (58.45) thousand dollars in mean (median) on-chain transactions. The standard deviation of 61.07 suggests that some days see extremely large (small) volumes, possibly motivated by news or other events. The skewness (2.03) and kurtosis (6.70) show that the distribution is right tailed and not too different from normal. The panel also reports the distribution quintiles. Twenty percent of the transfers are above 98,850 (below 30,410) dollars, with a maximum daily transfer of 544,170 dollars. The quintile analysis provides further evidence of a right-tailed daily volume distribution.

3.3 TERRORIST EVENTS DATA

To create a list of terrorist events, we rely on the GTD, which provides comprehensive details on terrorist attacks. It is curated at the University of Maryland and has been used in several academic papers (e.g., Cuculiza et al. [2021]).³¹ The data set contains 21,323 unique events from 2015 to mid-2019 in multiple countries. We collect information on the number of dead and wounded, location, method, and other details, such as whether a ransom was paid and the total property damage for each event in the data set. These details are essential for filtering the collected data of terrorist attacks for attacks in which financing is potentially in place. For instance, one may consider the total number of victims, the estimated property damage, and the number of operatives on the ground as proxies for the scale of the event.

As the choice of an attack's scale may be subjective and prone to sampling errors, leading to selection concerns in the empirical analysis, we rely on a more agnostic approach guided by global media coverage. A highly publicized and visible terrorist event is likely to have many victims, which involves complex operations (e.g., multiple bombings) and several operatives on the ground. To identify terrorist attacks that received global media

³¹ More information is available at <https://start.umd.edu/gtd/>.

attention, we check whether a standalone Wikipedia article exists for every event in the GTD sample. To have a standalone Wikipedia article, a terrorist attack must be (1) notable, (2) caused by violent nonstate actors, (3) unrelated to drug wars or cartel violence, (4) unrelated to ongoing military conflicts, and (5) in compliance with guidelines for the definition of terrorism.³² The first item is particularly relevant for filtering the collected list of terrorist events because notable attacks are largely documented by the press and media. For instance, the Sri Lanka Easter bombing has a standalone Wikipedia article based on 351 sources (including *BBC News* and *The New York Times*). This selection process results in 327 notable attacks.³³ We use these notable attacks as our main specification and the initial full sample of attacks as our baseline specification.

Table 2 details the sample composition. Panel A reports the filtering criteria, the number of events (multiple events per attack), the number of unique events (terrorist attacks), and the average number of victims (dead and wounded) in each sample. In line with our focus on large-scale events, the average number of victims is approximately five times larger in a notable attack (26) than in the baseline specification (5). Panel B reports additional attack characteristics. Notable attacks are usually carried out with bombs (77.47%) and occur mostly in the Middle East, North Africa, and South Asia (75.88%). In approximately one-third of the attacks, the property damage is less than one million dollars.

The sample of notable attacks is dominated by regions where local terrorist groups (e.g., ISIS, al Qaeda, al-Shabbab, and Boko Haram) operate. As local operations are likely to use more practical methods than Bitcoin (e.g., Dion-Schwarz et al. [2019]), we form an additional sample from which we exclude locations where the attacks are mostly carried out by local groups or political insurgents. The resulting data set of foreign attacks contains 50 unique events, mostly carried out in Western Europe (39.02%) and North America (29.27%). Bitcoin is more likely to finance activities distant from the home territory of the terrorist group (e.g., to circumvent capital controls and anti-terrorist financing measures); therefore, the refined vector of notable foreign attacks is more likely to contain relevant attacks for the purposes of this study. We use foreign terrorist attacks in the predictive part of the paper (section 7.1) to improve the model's accuracy.

4. Main Analysis: Abnormal Volume Around Terrorist Attacks

4.1 PREDICTIONS

Our goal is to examine whether outsiders can exploit blockchain's transaction transparency and identify flows of funds associated with terrorist attacks. We argue that large-scale, highly visible (notable) attacks need

³² More information is available at https://en.wikipedia.org/wiki/List_of_terrorist_incidents.

³³ See the online appendix for more details on sample construction.

TABLE 2
Terrorist Attacks Sample Description

Panel A: Summary statistics			
	Baseline	Main	Foreign
Sample criteria	[1]	[1–4]	[1–5]
Number of total events	53,321	1,012	82
Number of total unique events (<i>N</i>)	21,323	327	50
Number of dead (mean)	2.62	10.22	7.85
Number of victims (mean)	5.13	26.07	27.89
Panel B: Terrorist attacks characteristics (percentage)			
	Baseline	Main	Foreign
Type:			
Armed assault	22.35	21.54	47.56
Bombing/explosion	46.75	77.47	47.56
Hijacking	0.35	0.98	4.88
Others	30.55	–	–
Region:			
Eastern Europe	1.92	1.19	6.10
Middle East and North Africa	38.34	47.13	4.88
North America	0.75	2.37	29.27
South Asia	29.81	28.75	9.76
Southeast Asia	8.12	6.62	1.22
Sub-Saharan Africa	16.8	9.58	–
Western Europe	2.21	3.16	39.02
Others	2.06	1.19	9.76
Others:			
Ransom	0.56	–	–
Property damage (>\$1 M)	0.04	0.19	–
Property damage (<\$1 M)	24.44	28.56	25.61

This table reports descriptive statistics on the composition of the terrorist attacks sample. Panel A reports the number of terrorist events in each of the three samples used in the analyses together with the distribution of the number of dead and victims (i.e., wounded and dead). Panel B reports additional statistics for the samples used in the paper. We use the following sample selection criteria: [1] Sample of all GDT events in the period from January 2015 to June 2019 (retrieved on December 7, 2021). [2] Exclude attack types: Assassination, Hostage Taking, Infrastructure Attacks, Unarmed Assaults and Unknown. [3] Exclude attack method: Arson (Incendiary), Knife and Melee. [4] Notable events: events matched to standalone Wikipedia page. [5] Exclude the following locations: Afghanistan, Bangladesh, Burkina Faso, Cameroon, Central African Republic, Colombia, Democratic Republic of the Congo, Egypt, India, Indonesia, Iraq, Israel, Ivory Coast, Jordan, Kenya, Kuwait, Lebanon, Libya, Mali, Myanmar, Nepal, Niger, Nigeria, Pakistan, Philippines, Saudi Arabia, Somalia, South Sudan, Sudan, Syria, Thailand, Tunisia, Turkey, Ukraine, Venezuela, West Bank and Gaza Strip, Yemen, and Zimbabwe. See the online appendix for more details.

financing to buy weapons on the black market and pay terrorists on the ground. If terrorist financiers are worried about eliminating traces in the blockchain, they may use money-laundering techniques. If these techniques involve repeated reshuffling in several wallets to reduce traceability

of funds, then laundering creates excess volume.³⁴ We therefore expect an abnormally large volume to appear before an event.

Although we expect the volume after the attack to return to its normal level, there may be reasons for the volume to decline even further. Notable terrorist attacks are negative news, potentially suppressing Bitcoin trading. The suppression could also reflect a reduction in activity in unrelated illicit transactions that seek to fly under the radar of agencies searching for terrorist linkages due to the attack. Nevertheless, other reasons predict that volume could remain abnormally high after the event. For instance, if terrorist financiers make Bitcoin transfers to pay the perpetrators' families after the event using Bitcoin (as compensation for terrorist deaths), we may observe a higher-than-normal volume (Levitt and Venkatesh [2000]). Thus, Bitcoin volume patterns in the weeks after a large-scale terrorist attack are difficult to predict, which is why we focus our search for suspicious volume patterns to the weeks beforehand.

Additionally, abnormal volume is expected to increase in terrorist financing needs. 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. Notable events, those that lead to many injured and dead, are usually the consequence of a more significant level of logistics (e.g., several perpetrators) and equipment (e.g., mass shooting); therefore, they demand more funds to finance operations.³⁵

Moreover, we expect that Bitcoin transfers associated with terrorist attacks are likely to go through *Exchange* and *Mixer* service providers. While the former provides users with the option to cash out or to use an exchange wallet for payments and transfers, the latter is used to make Bitcoin tracing more difficult. We assume that whereas terrorist organizations use professional mixers or mixer-like services to reshuffle Bitcoins, terrorists on the ground use exchanges for payments and withdrawals (figure 1).

Finally, a significant variation in the regulation of crypto exchanges coupled with flawed compliance measures suggests that terrorist funds are likely to flow mostly through unregulated exchanges. For instance, unregulated exchanges often do not observe compliance procedures—such as KYC and AML—and are more likely to engage in deceiving activities to attract demand (Aloosh and Li [2021], Amiram, Lyandres, and Rabetti [2021], and Cong et al. [2021]).

³⁴ See the online appendix for a detailed network analysis of al Qassam wallets. The reshuffling factor is 38 laundering transfers to one final transfer or a laundering volume of 161,386 dollars to a total transfer amount of 4,246 dollars.

³⁵ The number of injured and dead is a noisy proxy for the costs of the attack. There could be a costly attack that ultimately 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.

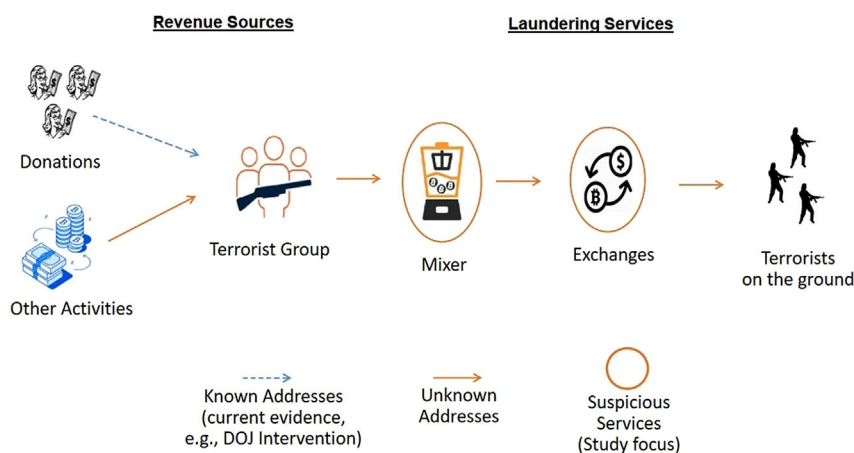


FIG 1.—Mechanics of terrorist financing. This figure plots the mechanics of terrorist financing. The left-hand side (Revenue Sources) illustrates the terrorist group’s potential sources of income from donation campaigns (known addresses) and other activities (unknown addresses). Most of the current evidence on terrorist financing (e.g., DOJ intervention in donation campaigns) is obtained from known addresses. Thus, the activities by unknown addresses (e.g., financing terrorist attacks) are not accounted for as part of the total Bitcoin volume associated with illicit activities. The right-hand side of the plot (Laundering Services) illustrates how terrorist groups may use blockchain-based currencies and services to finance large-scale terrorist attacks. Although the addresses of terrorists on the ground are unknown, by focusing on the services associated with camouflaging the transfers and funding the activities, outsiders can analyze blockchain data to infer users from transfers and interactions—the focus of our study.

5. Methodology

The accounting literature uses 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 et al. [2006]), the role of the media in disseminating news (Rogers, Skinner, and Zechman [2016]), and investors’ reaction to blockchain-related disclosures (Cheng et al. [2019]). In line with these abnormal volume analyses, we use an event study design that examines Bitcoin volume in the vicinity of large-scale attacks. If a terrorist attack is financed with Bitcoin and money-laundering techniques are used 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.³⁶

³⁶ See the online appendix for an example of mixed transactions between Bitcoin donors and the terrorist group al Qassam Brigades.

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 the event window. The event window comprises the two weeks before (after) the event. We then mean-adjust the realized volume in the event window by subtracting the mean volume from the benchmark window, both in logarithmic terms (Beaver [1968], Copeland [1979], and Bamber [1987]),

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

where abnormal mean-adjusted volume AV is calculated at 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 group N_t . The cumulative mean-adjusted abnormal volume CAV is constructed by the sum of AAV in the specified event windows T .

$$AAV_t = \frac{1}{N_t} \sum_{u=1}^N AV_{u,t},$$

$$CAV_t = \sum_{t=1}^T AAV_t; \quad t = 1, \dots, n.$$

We report the mean CAV for the period before $[t - 15, t - 1]$ and after $[t + 1, t + 15]$ the event to check for abnormality. We base the choice of the event window on two considerations. First, there are two types of cryptocurrency transfers associated with terrorism. The first type is donation campaigns, which happen over the long term and between donors and the terrorist organization. The other type is direct financing between the terrorist organization and terrorists on the ground. The latter has a short-term characteristic (e.g., a machine gun purchase a week before the Paris attack) and is the focus of our analysis. Second, even though we focus on notable events, minor terrorist attacks occur very frequently. Therefore, we keep our event and estimation windows relatively short so that the total window $[t - 35, t + 15]$ is less likely to be contaminated by any other similar event.

As an additional step to deal with multiple terrorist attacks clustered in calendar time, we report all CAV responses as the mean CAV results from bootstrapped analyses. The bootstrapping randomly selects 10 events from the vector of terrorist attacks in each run. The procedure is carried out 50 times, and the mean- CAV results across all bootstrapped analyses are reported in the tables. This method is important for two reasons. First, as we are unsure which event within the vector of terrorist attacks is financed with Bitcoin, bootstrapping mitigates this potential selection concern. Second, the method also mitigates the influence of outliers and other factors unrelated to terrorist financing that may confound the results around some specific dates.

The test statistic in the event study could be affected by two issues. The first is a potential cross-sectional correlation of abnormal volume due to shared event dates across users.³⁷ The second is that if unobserved events cluster in time with observed events, this can lead to event-induced volatility (Bernard [1987]). Consequently, both issues potentially introduce downward bias in the standard deviation, leading to an overrejection of the null hypothesis. To address these concerns, we use a nonparametric test, suggested by Corrado and Zivney [1992], that corrects for induced volatility and cross-correlation (Campbell and Wesley [1993]). One disadvantage of this adjustment is that the test might lose power for larger event windows (e.g., over 30 days).

We follow the Corrado [1989] rank test and transform abnormal volumes into ranks in both the event and benchmark periods. If ranks are tied, the midrank is used. To adjust for missing values, Corrado and Zivney [1992] suggest a standardization of the ranks by the number of nonmissing values plus 1, as follows:

$$\bar{K}_{u,t} = \frac{\text{rank}(AV_{u,t})}{M_u + L_u + 1},$$

where M_u (L_u) refers to the number of nonmissing values in the (benchmark) event window. For multiple events, Campbell and Wesley [1993] consider the sum of the mean excess rank for the event window,

$$\begin{aligned}\bar{K}_{T_1, T_2} &= \frac{1}{L_2} \sum_{t=T_1+1}^{T_2} \bar{K}_{u,t}, \\ S_{\bar{K}}^2 &= \frac{1}{L_1 + L_2} \sum_{t=T_0}^{T_2} \frac{N_t}{N} \left(\bar{K}_t - \frac{1}{2} \right)^2,\end{aligned}$$

where $L_1 = T_1 - T_0 + 1$ denotes the benchmark window length, with T_0 (T_1) representing the earliest (latest) day of the benchmark window; furthermore, $L_2 = T_2 - T_1$ denotes the event window length, with T_1 (T_2) denoting the earliest (latest) day of the event window. Therefore, we obtain (H_0 : $CAV = 0$):

$$t_{\text{stat}} = \sqrt{L_2} \left(\frac{(\bar{K}_{T_1, T_2} - \frac{1}{2})^2}{S_{\bar{K}}} \right).$$

For comparison, we also compute standard tests across all model specifications (not tabulated). However, the test statistics of the standard approach are often larger than those of the nonparametric approach (Corrado and Zivney [1992] and Campbell and Wesley [1993]). Therefore, we opt for the more conservative nonparametric test statistic to prevent reporting significance tests with potential downward biased standard errors.

³⁷Brown and Warner [1985] show that cross-sectional tests are prone to event-induced volatility.

TABLE 3
Abnormal Volume Analyses for All Users

Panel A: Terrorist attacks and other impactful events					
Period	Main ^a	Baseline ^b	Holidays	Peaks	Hardforks
Days before	9.53 (2.73)	2.73 (0.29)	−29.74 (−3.35)	39.78 (12.73)	83.64 (1.42)
Days after	−7.33 (−1.30)	1.90 (2.29)	−2.34 (−0.71)	−17.79 (−9.74)	−22.16 (−9.10)
Panel B: Expanded event window (main sample)					
Event Window	15 Days ^a		30 Days ^a		45 Days ^a
Days before	9.53 (2.73)		12.51 (2.14)		14.10 (0.54)
Days after	−7.33 (−1.30)		−6.87 (−1.96)		−6.92 (−2.58)
Panel C: Expanded benchmark window (main sample)					
Benchmark Window	20 Days ^a		90 Days ^a		180 Days ^a
Days Before	9.53 (2.73)		12.46 (1.74)		20.15 (0.89)
Days After	−7.33 (−1.30)		−6.61 (−1.84)		−18.29 (−1.88)

This table reports the cumulative mean-adjusted abnormal volume (CAV) surrounding impactful trading events and terrorist attacks (panel A), for varying event windows (panel B) and for varying benchmark windows (panel C) surrounding the terrorist attacks in the main sample for all Bitcoin users ($U = 342$). The regular event window consists of the days in the intervals $[-15, -1]$ and $[1, 15]$, respectively. Volume is calculated as the logarithmic change of the sum of inbound and outbound volumes. Unless indicated otherwise, abnormal volume is mean adjusted for estimates generated in the period of 20 days before the first day in the event window (regular benchmark window). The vector of terrorist attacks is selected through bootstrapping in a total of 50 runs. Each run randomly selects ten unique events from the respective sample and excludes events coinciding with impactful events. CAV responses are then averaged across all bootstrapped vectors, and the mean-adjusted CAV response is reported in percentage. The t -statistic is reported in parentheses (see subsection 4.2 for details on the computation of the t -statistic). ^adenotes the main specification based on notable attacks ($N = 327$). ^bdenotes the baseline specification based on all attacks, including minor ones ($N = 21,323$).

6. Results

The left-hand side of table 3, panel A, reports the results for the main specification (notable terrorist events) and the baseline specification (all terrorist events). For each analysis, the bootstrapped mean-CAV response is adjusted for impactful events.³⁸ The main specification mean-CAV response is 9.53% (−7.33%) in the weeks before (after) the event. The increase in the mean-CAV response before the event is consistent with the

³⁸The concern that large-scale terrorist attacks occur at a higher rate during holidays, as terrorists aim to maximize the number of deaths, or during other events such as hardforks and price peaks is mitigated by excluding terrorist attacks that fall in the vicinity of these impactful events in each bootstrap run.

increase in money laundering to satisfy unexpected demand from terrorist financiers. The decrease in the mean CAV response after the event is statistically insignificant. The baseline specification mean-CAV response is 2.73% (1.90%) in the weeks before (after) the event. The coefficient magnitude and the statistical significance decrease in the baseline specification, consistent with minor terrorist attacks needing little or no immediate financing. In other words, abnormal Bitcoin volume (as a proxy for terrorist financing) increases with the estimated scale of an attack.

To provide context, we also calculate CAV responses to significant events unrelated to terrorist attacks that are likely to impact Bitcoin. As opposed to events used in the accounting literature, such as earnings announcements and stock splits, the literature on cryptocurrencies is still in its infancy. Therefore, by analyzing impactful events, we provide a context for the interpretation of our main results. To capture CAV responses to impactful events, we create vectors of events, such as holidays, price peaks, and hardforks. *Holidays* is a vector of dates for the main western holidays, such as Christmas and Easter. *Peaks* is a vector of dates for when the Bitcoin price breaks the resistance price. *Hardforks* is a vector of dates for Bitcoin blockchain splits.³⁹

The right-hand side of panel A reports the results for benchmark events. CAV responses to holidays are -29.74% (-2.34%) in the weeks before (after) the event. This result suggests a reduction in Bitcoin activities around holidays such as Christmas and Easter. CAV responses around (Bitcoin) price peaks are 39.78% (17.79%) in the weeks before (after) the event. This result is consistent with increasing Bitcoin activity once it is near its price resistance. CAV responses around Bitcoin hardforks are 83.64% (-22.16%) in the weeks before (after) the event. Hardforks are potentially the most impactful events on Bitcoin volume, as they work as a dividend for existing Bitcoin holders (i.e., holders receive equivalent assets in the newly created blockchain). Although the mean-CAV responses for impactful events are several orders of magnitude larger than the mean-CAV responses for terrorist attacks, the results suggest that blockchain data are a significant source of information to predict other sizeable events.

To test the robustness of our findings, we estimate the mean-CAV response across different model specifications. In panel B, we expand the event window to $[t - 30, t + 30]$ and to $[t - 45, t + 45]$. The mean-CAV response for 30-day event windows resembles the main specification. However, when we expand it to 45 days, the mean-CAV response increases for the periods before the event. This result suggests induced volatility in the event window due to overlapping events, which could lead to overestimated coefficients. In panel C, we expand the benchmark windows to 90 and 180 days, respectively. Expanding the benchmark window, which incorporates other similar events, has a similar effect of inducing volatility in the

³⁹ See the online appendix for details on these events.

TABLE 4
Abnormal Volume Analyses by Business Groups

Panel A: Exchanges				
Period	All (U = 101)	Regulated (U = 10)	Unregulated (U = 91)	Diff (U = 101)
Days before	9.28 (5.18)	4.11 (2.24)	12.72 (5.47)	8.61 (3.97)
Days after	-5.26 (-3.32)	-9.07 (-0.90)	-8.30 (-3.18)	0.78 (0.19)
Panel B: Other services				
Period	Mixer (U = 36)	Gambling (U = 50)	Services (U = 56)	Dark Markets (U = 99)
Days before	11.09 (3.85)	-4.05 (-2.36)	-9.76 (-3.16)	-7.82 (-3.04)
Days after	-7.00 (-3.86)	1.59 (4.47)	4.86 (5.01)	2.51 (5.33)

This table reports the cumulative mean-adjusted abnormal volume (CAV) for crypto exchanges (panel A) and other business groups (panel B). The number of users in each group is reported following “U =”. The event window consists of the days in the intervals [-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 vector of terrorist attacks is selected through bootstrapping in a total of 50 runs. Each run randomly selects ten unique events from the respective sample and excludes events coinciding with an impactful event. CAV responses are then averaged across all bootstrapped vectors, and the mean-adjusted CAV response is reported in percentage. The *t*-statistic is reported in parentheses. We use the terrorist attacks from the main specification in this table ($N = 327$).

mean-CAV response, but the statistical power dissipates as the window increases. We conclude that due to events clustering in time and the potential contamination of the results by longer windows, the narrower window we use seems adequate.

Our next set of tests examines which type of Bitcoin service is more likely to be associated with terrorist attacks. Table 4, panel A, reports the mean-CAV response for *Exchange* services. The mean CAV response for all exchanges is 9.28% (-5.26%) in the weeks before (after) the event. This result supports our expectation that *Exchange* channels concentrate most of the illicit activity, as operatives on the ground need to cash out funds to use them for payments of local operations. We then separate the effects for regulated and unregulated exchanges.⁴⁰ We do this for two reasons. First, regulated exchanges are less likely to have flows of illicit funds than unregulated exchanges due to regulatory oversight. For instance, some regulated exchanges, such as Coinbase and Kraken, must comply with NYSDFS (New York State Department of Financial Services) guidance, keep a relatively high level of compliance measures in place (such as KYC and AML

⁴⁰ Regulated exchanges are services licensed by financial regulatory authorities in the United States, United Kingdom, or Japan.

rules) and are subject to further scrutiny by other U.S. regulators (i.e., the Internal Revenue Service and the Securities and Exchange Commission). Second, the separation permits us to provide insights into the ongoing debate over cryptocurrency regulation (e.g., Amiram et al. [2021], Cong et al. [2022], Griffin and Shams [2020], Foley et al. [2019], Fusaro and Hougan [2019], and Sokolov [2021]). The difference in the mean-CAV responses for unregulated and regulated exchanges is 8.61% (0.78%) in the weeks before (after) the event. These results indicate that most of the funds likely to finance terrorist attacks go through unregulated exchanges' wallets.

Panel B reports the mean CAV response for the remaining services. The response for Mixer is significantly larger in the weeks before terrorist attacks. This pattern is consistent with unexpectedly large volumes of Bitcoin being reshuffled in the period before an attack. The result supports our prediction that mixers (together with the unregulated exchanges discussed above) are potential channels for terrorist funding. Additionally, the cross-sectional variation, such as the increased use of mixers, corroborates our previous robustness test that confounding events are unlikely to drive the effects we document. Moreover, as mixers and other blockchain services are highly interconnected (figure 2), a large, unexpected demand for mixers likely crowds out other services in the chain. During normal times, mixer activities are mainly used to camouflage transfers between users and illicit services (Foley et al. [2019]).⁴¹ As mixer volume increases in the vicinity of terrorist attacks, increased fees and slower reshuffling services would then temporarily crowd out other less time-sensitive users. The mean CAV responses to *Gambling*, *Services*, and *Dark markets* are negative (positive) 4.05% (1.59%), 9.76% (4.86%), and 7.82% (2.51%), respectively, in the weeks before (after) a terrorist attack occurs. While the negative responses before the event suggest crowding out, the positive responses in the weeks afterward suggest a normalization of these services to the users. These patterns further mitigate concerns that earlier results are mechanical.

We next distinguish between terrorist attacks carried out locally and in foreign territories. Foreign attacks are mostly perpetrated by ISIS and al Qaeda, organizations known for having acquired both the resources and the expertise in cryptocurrency usage to finance large-scale operations. Table 5, panel A, reports the bootstrapped mean CAV results for foreign and local events. The results suggest that CAV responses are larger for terrorist attacks occurring in foreign versus local territories, consistent with cryptocurrencies being used to evade capital restrictions and anti-terrorist financing measures in traditional cross-country money-transfer venues. To illustrate the effects of individual events composing the vector of foreign attacks, we plot the CAV responses for three large events in panel B: Charlie Hebdo, an al Qaeda-claimed attack that occurred in January 2015; Manchester Arena, an ISIS-claimed attack that occurred in May 2017; and the Sri

⁴¹ Some exchanges, such as Coinbase, ban users that transfer coins directly to dark market services. Therefore, users tend to use mixer services to camouflage the link.

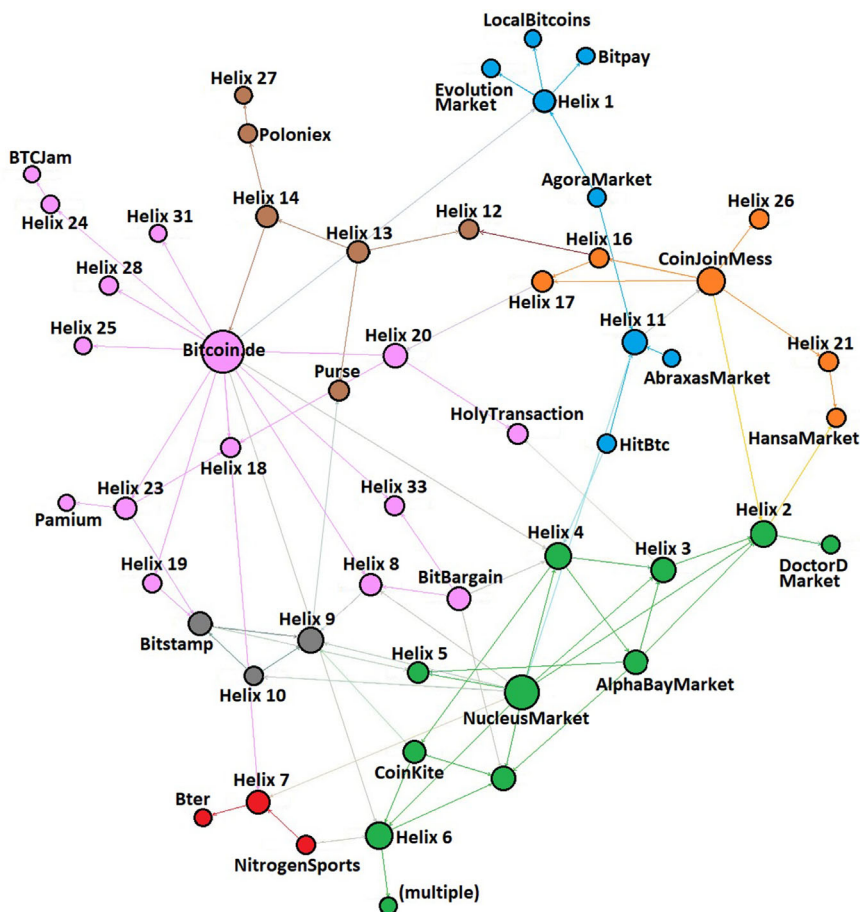


FIG 2.—Blockchain services network analysis. This figure illustrates the interconnectedness of a sample of mixer services (Helix) with other blockchain services. The figure shows that mixer services are commonly used by other Bitcoin users, such as crypto exchanges (e.g., Bitstamp, LocalBitcoins, and Poloniex), gambling (e.g., NitrogenSports), dark markets (e.g., AlphaBayMarket, AbraxasMarket, and EvolutionMarket), and other services (e.g., BitPay, CoinJoinMess, and CoinKite). The figure was created with 52 nodes and 85 edges. The range was reduced to improve readability. The size of each node measures centrality (i.e., larger nodes are central to several transactions), and colors show proximity (i.e., nodes near each other are in the same color).

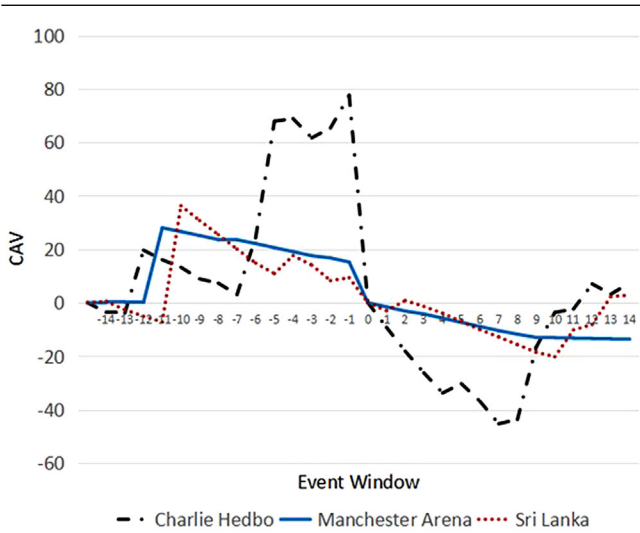
Lanka Easter bombing, an ISIS-claimed attack that killed more than 250 people and wounded hundreds in April 2019.⁴²

⁴²The GTD database underreports the number of victims in this event. We consider the number of victims according to *The New York Times* article: “Sri Lanka Attacks: What We Know and Don’t Know,” April 24, 2019 (<https://www.nytimes.com/2019/04/24/world/asia/sri-lanka-easter-bombing-attacks.html>).

TABLE 5
Abnormal Volume Analyses by Location and Type of Event

Panel A: Attacks by location and type (CAV)						
Period	Notable Attacks		All Attacks by Type			
	Foreign ^a (N = 50)	Local ^b (N = 277)	Ransom (N = 289)	Damage (N = 8,919)	Melee (N = 1,181)	Arson (N = 2,154)
Days before	28.39 (3.22)	2.71 (0.59)	3.57 (1.79)	1.19 (0.12)	1.69 (0.89)	1.53 (0.93)
Days after	-4.59 (-1.89)	-3.56 (-2.07)	4.18 (2.89)	1.80 (0.21)	1.54 (0.74)	0.87 (0.39)

Panel B: Select foreign attacks



This table reports the cumulative mean-adjusted abnormal volume (CAV) for various types of terrorist attacks (panel A) and plots the CAV around select foreign terrorist attacks (panel B) for all Bitcoin users ($U = 342$). The event window consists of the days in the intervals $[-15, -1]$ and $[1, 15]$. 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 vector of terrorist attacks is selected through bootstrapping in a total of 50 runs. Each run randomly selects ten unique events from the respective sample and excludes events coinciding with an impactful event. CAV responses are then averaged across all bootstrapped vectors, and the mean-adjusted CAV response is reported in percentage. The t -statistic is reported in parentheses. ^adenotes the main specification but is limited to foreign events. ^bdenotes the main specification but is limited to local events. The following filters do not include notable attacks: Ransom, a vector of terrorist attacks for which information on ransom payment is available; Damage, a vector of terrorist attacks for which information on property damage is available; Melee, a vector of melee-based terrorist attacks; and Arson, a vector of arson-based terrorist attacks.

Finally, we carry out a series of additional robustness tests. One concern is that the documented negative CAV responses in the weeks after the event are spurious (e.g., driven by the research design or the model specification). This concern is plausible because negative CAV responses are also observed for impactful events such as price peaks and hardforks. A potential explanation is that the CAV increases with transaction costs and

therefore suppresses volume in the post-event window. Consistent with this argument, Lennart [2020] finds that transactions with high levels of information asymmetry negatively affect abnormal trading volume once the event becomes public knowledge. To mitigate this concern, we test a vector of terrorist attacks for which ransom payments have been made. Because ransom payments are made after the event happens, we expect positive mean-CAV responses in the post-event window. Consistently, the mean CAV response in the post-event window is 4.18% for a sample of ransom paid attacks.

Furthermore, we also test samples based on whether the information on property damage is available or the attack was carried out using a low-tech method (e.g., a melee or arson attack). While the former provides a randomly chosen sample—because information on property damage is limited—the latter provide subsamples of attacks unlikely in need of financing.⁴³ Therefore, we expect no substantial effects in any of these subsample tests. The results support our assumptions, mitigating concerns about the research design and the sample selection.

7. Case Study: The Sri Lanka Easter Bombing

The event-study methodology we used in the previous section has two primary limitations. First, parties interested in researching or tracking 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 terrorism. We address these limitations by comprehensively analyzing anomalous transfers in the vicinity of the Sri Lanka Easter bombing. This case study-type analysis also provides further insights into our model for terrorist attack prediction. The motivation to examine the Sri Lanka Easter bombing stems from rumors that arose after the event claiming that Bitcoin had been used to finance it. However, there is no information about the users and wallets involved, so we employ several forensic accounting techniques to pinpoint anomalous transfers at the individual-user level and to track the funds on the blockchain.⁴⁴

We follow Laptev, Amizadeh, and Flint [2015] and Dai et al. [2019] and employ a rolling three-sigma rule to detect anomalous transfers among individual users in the vicinity of the bombing. We consider a transfer anomalous if the amount exceeds three standard deviations of the user's last three

⁴³ Alternatively, we run several placebo-tests based on randomly selected dates. However, because one may argue that terrorist attack dates were not randomized, we rely on chosen filters (such as property damage information) whose incomplete information induces randomization. For instance, many attacks without information on property damage likely resulted in property damage.

⁴⁴ We are thankful to Whitestream for providing additional evidence of terrorist-associated transfers for this event and to DeepSeek for discussions on the mechanics of terrorist operations.

TABLE 6
Anomalous Transfers Surrounding Sri Lanka Easter Bombing

Groups	I Users	II >3σ	III #	IV #/User	V Mean	VI Max
Dark Markets	3	0.02	11	3.67	57.86	255.48
Exchange	25	0.02	67	2.68	2,163.24	54,320.28
Gambling	8	0.02	22	2.75	18.15	194.00
Mixer	3	0.01	7	2.33	7,581.20	26,122.56
Services	9	0.02	28	3.11	188.84	1,382.84
Total	48	0.02	135	2.81	1,642.56	54,320.28

The table reports three-sigma rule anomaly detection results for 48 flagged users whose anomalies occur during April 2019. Column I shows the number of flagged users per respective business group. Column II reports the number of transfers (in percentage of total) larger than three standard deviations. As a benchmark, three standard deviations fall in 0.27% of a normal distribution. Column III reports the number of flagged transfers per user, and column IV shows the average flagged transfer per user. Columns V and VI report the mean and maximum transfer amount in Bitcoin units, respectively.

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 could indicate the presence of an anomaly. Outliers, being the most extreme observations, may be suspicious if they appear around specific events—in our case, if they precede a large-scale terrorist attack.⁴⁵ The rule is well suited to our purposes because the transfer data tend to be normally distributed and large-scale terrorist attacks are rare.

Table 6 reports the results for the three-sigma rule test. The rule identifies 48 users (out of 342) with at least one anomalous transfer around the Sri Lanka Easter bombing (April 2019). On average, only 2% of the transfers of a user exceed three standard deviations—as a benchmark, only 0.3% of the values exceed three standard deviations in a normal distribution. *Dark markets* wallets have, on average, more anomalies than the other groups. The reason may be business-driven, as this type of business probably accumulates a certain number 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 markets* and *Exchange* wallets. *Mixer* anomalies are the largest, followed by *Exchange* anomalies. In early April, a large transfer initiated in a U.S.-based exchange had feedback effects on 38 users. As this transfer was market-based (the BTC price rose by 10% that day) and was made more than two weeks before the Sri Lanka Easter bombing, it is unlikely that these 38 users were associated with the event. Another 4 users’ anomalies are observed too far away from the event, that is, exceeding [−15,

⁴⁵ In probability theory, an event is considered to be practically impossible if it lies in the region of values of the normal distribution of a random variable at a distance from its mathematical expectation of more than three times the standard deviation (Pukelsheim [1994]). Laptev et al. [2015] use this approach to develop a generic and scalable system used by Yahoo to detect anomalies in large time-series data. Dai et al. [2019] use a similar approach to detect anomalies in accounting data.

+15]. The remaining 6 suspicious users are distributed into *Service* (3), *Exchange* (2), and *Dark markets* (1). The largest anomaly in the period 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 cryptocurrency, takes a fee, and sends fiat payments to businesses. By doing so, it mitigates exchange risks attributed to oscillations of cryptocurrency prices. The difficulty of accepting Bitcoin as payment is severe, as its volatility is much higher than that of widely used currencies (Yermack [2015]). Assessing the wallet's balance, we find that the mean balance increases significantly one day before the event and falls sharply on the day of the event, which is a pattern consistent with terrorist financing. A sharp decline is flagged at the event's end day at a value of 400 BTC (\$2.8 million).

We further investigate whether wallets associated with money laundering are used to feed the gateway's abnormally large transfers. We find a wallet constantly feeding transfers into the gateway. This wallet has over 1 million transactions and approximately 1.9 million worth of 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 find 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.⁴⁶

We next exploit the fact that some of the funds were transferred to a crypto exchange. These funds likely were (1) exchanged to fiat and withdrawn, (2) sent to several other addresses, or (3) exchanged for another cryptocurrency. Unfortunately, testing the first two options is impractical for two reasons. First, this exchange is one of the largest, containing thousands of daily transactions.⁴⁷ Second, exchange transfers may occur off the blockchain in the exchange platform. However, the third option can be tested by searching for an association between a particular user, date, and value in other blockchains. To do so, we searched for the gateway's associated wallets in other cryptocurrencies, such as Litecoin, Ripple, and Bitcoin Cash. A few wallets denominated in these cryptocurrencies were found. However, a wallet in Ripple stands out as the best candidate because the timing, amount, and direction of the transfers match the information we have on the exchange. The account was activated in December 2014 and

⁴⁶ An address can be verified for reported crimes, such as extortion and ransomware, through services such as <https://bitcoinwhoswho.com/> and <https://www.bitcoinabuse.com/>.

⁴⁷ Although Bitcoin transactions are traceable, Bitcoin units are not. For instance, if two Bitcoins are transferred from wallet A to wallet B already containing one Bitcoin, and then another two Bitcoins are transferred from wallet B to wallet C, we know that one Bitcoin in C came from A but we are not certain about the origin of the second Bitcoin.

TABLE 7
Analysis of Ripple Transactions for the Gateway Ripple Wallet

Panel A: Ripple transfers (January to April 2019)		
	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 (January to April 2019)		
Bins	Observations	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%
Panel C: Volume frequency (before 2019)		
Bins	Observations	Frequency
Below 1,000	217,114	98.587%
1,001 to 10,000	2,990	1.358%
10,001 to 100,000	121	0.055%
100,001 to 500,000	0	0%
Above 500,000	0	0%

The table reports descriptive statistics for Ripple transfers and volume. Panel A reports inflows and outflows at the Gateway Ripple wallet for the period of January to April 2019. Flows are aggregated daily and measured in thousands of XRP (Ripple currency). We then report volume frequencies for the Gateway Ripple wallet for the period of January to April 2019 (panel B) and for the period before 2019 (panel C). Frequencies are based on five bins of XRP values. Consistent with the business model (small payments), transfers below 1,000 XRP are the most frequently occurring and represent 90% of transfers in the period. However, a few extremely large transfers occurred during the period from January to April 2019. In contrast, this user made no transfers above 100,000 XRP prior to 2019.

has over 250,000 transactions and a balance of over 17 million XRP (\$5.15 million) at the time of this writing.

Table 7, panel A, shows the daily volume of transactions at the gateway Ripple wallet (GRW), which we flagged for suspicious behavior during the first half of 2019. The wallet moved, on average, 133 (153) thousand XRP in daily inbound (outbound) transfers. However, April 23 is an atypical day during that period, as seen in figure 3. On typical days, the GRW averages 200,000 XRP in total transfers. However, most of these transfers involve less than 1,000 XRP. A large inbound transfer of 660,000 XRP (\$212,000) arrived on the morning of April 23, stayed in the GWR wallet for approximately 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




FIG 3.—Historical Ripple transfers. The plot shows aggregated daily inbound (blue), out-bound (red) and total (black) transfers on the Ripple network at the Gateway Ripple wallet (GRW) in 2019. XRP transfers are measured in units. An abnormally large transfer (0.7 million XRP) occurred in the early morning one day after the Sri Lanka Easter bombing.

Bitcoin blockchain. Second, the XRP transfer comes from the same exchange where BTC funds were sent. Finally, these transfers are outliers relative to the average historical transfer size in the GRW. To understand how abnormal these transfers are, as seen in 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 collecting crypto payments from online stores. Larger transfers may relate to vault storage or other services and are not obvious red flags, but extremely large transfers (>100,000 XRP) are inconsistent with the gateway’s business model. As a benchmark for the changes in the transfer behavior of this particular user, panel C shows that before the period likely associated with the financing of the Sri Lanka Easter bombing, the vast majority of transfers were below 1,000 XRP (consistent with the business model), very few transfers were more than 10,000 XRP, and no transfers exceeded 100,000 XRP.

The flagged anomalous transfers allow us to track them a step further on the Ripple network. Analyzing the network of associated wallets, we flag three other wallets as having a peculiar behavior that resembles money laundering. The first wallet behaves as a mixer, reshuffling large sums 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 August 2021, the third wallet had a balance of over \$200 million, making it one of the wealthiest wallets on the Ripple blockchain.⁴⁸

To conclude, the evidence in this section suggests that a chain of sophisticated money-laundering wallets in both the Bitcoin and Ripple blockchains was associated with anomalous transfers in the vicinity of the Sri Lanka Easter bombing. As the data behavior, such as the type, amount, and timing of transfers and wallets associated with the terrorist attack, is a potentially rich source for predicting terrorist attacks, we exploit this information in the next section.

7.1 PREDICTIVE ABILITY OF BLOCKCHAIN DATA

In this section, we test whether we can use the above insights to construct a model that increases our ability to predict terrorist attacks. To do so, we employ three different machine-learning models in the training set and choose the best performer for the analysis in the validation set. We consider SVM, NN, and RF machine learning models, as they stand out as state-of-the-art solutions of supervised nonlinear learning classifiers. Although we must calibrate the SVM to provide reasonable results, the NN and RF are specification-free. A drawback of NN is that the parameters are harder to interpret compared to the other two approaches. Despite these differences between models, we cannot anticipate which model best fits our purposes.

We identify the model with the best fit by using similar tuning (model calibration) across models and choosing the one that has the best performance. Several performance metrics can be used, as we elaborate below.⁴⁹ We train the models using the flagged user, discussed in the previous section, that had several anomalous transfers in the vicinity of the Sri Lanka Easter bombing. To define the parameters included in the classification model, we use predictors from the blockchain data that our case-study analysis has shown to be associated with terrorist attacks. For each model, we run the following formula:

$$\begin{aligned} \text{Error}(t) = & \text{Volume}_{(t-1)} + \text{In}_{(t-1)} + \text{Out}_{(t-1)} + \text{Life}_{(t-1)} \\ & + \text{Anonymous}_{(t-1)} + \text{Exchange}_{(t-1)} + \text{Mixer}_{(t-1)} \\ & + \text{DarkMarkets}_{(t-1)} + \text{Balance}_{(t-1)} + \text{Sigma}_{(t-1)} \quad (1) \end{aligned}$$

⁴⁸The wallet is among the top 0.1% richest wallets on the Ripple blockchain. See <https://ledger.exposed/richstats#percentage>.

⁴⁹We adjusted the performance of the SVM model using the Gaussian kernel function to approximate the NN and RF performances. As training classifiers in high dimensions (i.e., in the presence of several parameters or resampling) is time-consuming, the adjustment helps this model yield better results.

where *Terror* is an indicator variable that captures whether an attack occurred on that day. The list of terrorist events is the same as used in the main specification but is limited to foreign attacks. However, as several events are in the vicinity of impactful events, such as price peaks and hardforks, we exclude these impactful events from the training and validation samples to improve the machine-learning model's accuracy. This filtering shrinks the vector of foreign attacks from 50 to 30 events, mostly claimed by either ISIS or al Qaeda.⁵⁰ *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, that is associated with sending or receiving the funds from users in our sample. *Anonymous* is an indicator variable that captures whether funds were transferred 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 market address. *Balance* captures the total balance of the user's wallet in Bitcoin units. *Sigma* is an indicator variable that captures whether transfer size exceeds three standard deviations from 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.

The flagged user we employ for training purposes provides a set of 7,443,285 transfers in the period from 2015 to mid-2019. We split the set into training and validation sets, with the former (latter) having approximately 70% (30%) of the observations for the period. This results in 21 (9) events in the training (validation) sample. Rather than manually setting parameters that we believe have high predictive power, we tenfold cross-validate all three models (James et al. [2017]). The cross-validation combines averages of measures of fitness 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 each subset, and comparing the results to arrive at the model's best fit. Once the final tuning values are assigned, we refit the final model by using the entire training set. The technique also helps mitigate overfitting (Kuhn and Johnson [2018]). This process is repeated 100 times for each model.⁵¹

⁵⁰ Since the user data start from mid-April 2015, three events happening early in that year are not included in the analysis. We did not exclude holidays because of the user's association with the Sri Lanka Easter bombing. For robustness, we also report the training results for the original vector of notable events.

⁵¹ Although very unlikely, it is possible to implement resampling incorrectly. However, performing 100 repetitions in the resampling technique mitigates this issue.

Another concern arises from having too many binary predictors. As we tune models using resampling methods, a random sample of the training set may result in some binary predictors becoming a zero-variance predictor. 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-zero- 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 0.108. 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 of 0.108. 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 0.994.

Table 8, panel A, reports results across all models for the subsample of the training sample 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 there will be no terrorist attack. The high rate indicates that terrorist attacks are rare. However, all three models' accuracy suggests that the machine-learning models predict better than a naive guess. In other words, even for highly skewed no-terrorist attack occurrences, these models can predict attacks with better accuracy than a naive guess. However, as indicated by the lower Kappa, the SVM and NN results are not significant at the 95% confidence level. This 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 produce false positives (NO/YES), while the prediction (YES/NO) indicates that both models produce many false negatives. We follow McNemar [1947] and find no significant sampling error affecting the differences between correlated proportions in the confusion matrix table. Perhaps a more appropriate measure is *Balance accuracy*, which equally weights the accuracy for predicting positive and negative events. For instance, the naive predictor that there are never terrorist attacks has a balanced accuracy of 0.4968 $((0.994 + 0)/2)$. All models have a better balance accuracy than a naive guess. The tuning results suggest that the RF model stands out as the best fit for predicting terrorist attacks in the validation set.⁵² Table 8, panel B, reports the most important parameters in the RF model. *Life* appears to be the most important parameter, with a mean decrease accuracy (MDA) of 35.13%, followed

⁵² As a robustness test, we train the Random Forest model using the vector of notable events without adjusting for foreign attacks and impactful events. The results of this test are inferior to a naive guess, suggesting that Bitcoin financing is more likely to be used for attacks geographically distant from the home territory of the terrorist organization.

TABLE 8
Predicting Future Terrorist Attacks: Training Sample

Panel A: Training ($N = 21$ events)				
Model	SVM ^a	NN ^a	RF ^a	RF ^b
Accuracy–(ACC)	0.9939	0.9939	1	0.8830
Kappa	0.1076	0.0895	1	0.1685
No Information Rate–(NIR)	0.9937	0.9937	0.9937	0.9725
p -Value (ACC > NIR)	0.0525	0.0949	<0.0001	1
McNemar’s p -value	<0.0001	<0.0001	<0.0001	<0.0001
Balanced Accuracy	0.5292	0.5240	1	0.7222
Precision	0.9941	0.9940	1	0.8924
Recall	0.9999	0.9989	1	0.5514
F -measure	0.9969	0.9969	1	0.6816
Panel B: Parameter importance				
Parameter	0	1	MDA	MDG
Life	33.64	37.32	35.13	1,414.37
Balance	32.93	25.44	33.85	281.19
Volume	10.91	8.80	13.34	49.39
Anonymous	9.20	5.74	10.69	1.30
Exchange	7.76	1.83	6.50	0.99

This table reports machine learning results for the training sample. Panel A reports the results for the trained set across three machine learning algorithms. SVM stands for Super Vector Machine. NN stands for Neural Networks. RF stands for 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 terrorist attack occurred on that day. $Volume$ 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 the transfer size exceeds three standard deviations from the user’s historical mean over the last three months. All predictors used in the classification model are at the transfer-user level and lagged by one day. Panel B lists parameters’ importance in the tuned model in decreasing order of importance, where MDA stands for mean decrease accuracy, and MDG stands for mean decrease Gini. N indicates the number of unique events used in the training set. ^adenotes the main specification limited to foreign attacks after adjusting for *Peaks* and *Hardforks*. ^bdenotes the main specification (32 events in the training sample).

by *Balance* (33.85%), *Volume* (13.34%), *Anonymous* (10.69%), and *Exchange* (6.50%).⁵³ This result suggests that alternative models based on any single parameter, such as *Life*, can be exploited to improve predictability.⁵⁴

⁵³ Their relative importance measured as a mean-decreasing Gini (MDG) follows a similar order.

⁵⁴ However, there are two issues with an analysis that is based on the *Life* parameter. First, the proper identification of young wallets is challenging as these wallets could be created for the main purpose of illicit usage and, therefore, not be linked to known addresses. Second, Bitcoin mapping should occur in real time to reflect the vast number of newly created addresses.

TABLE 9
Predicting Future Terrorist Attacks: Validation Sample

Validation ($N = 9$ Events)				
	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	0.250	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

This table reports machine learning results for the validation sample. The first column reports statistics for the flagged user (i.e., GRW). The remaining columns report the mean statistics for each group of users. The results in this table are for a 10-fold cross-validated random forest model with 500 trees and nine variables at each split. N indicates the number of unique events used in the validation set. All specifications use the vector of notable events limited to foreign attacks after adjusting for *Peaks* and *Hardforks*.

We therefore apply our RF-trained model to the validation sample at the flagged user and to a group of users defined by their business type: *Exchange*, *Gambling*, and *Service*. The model accuracy improves in the group of users (table 9), suggesting that some of these users have transfers that follow the same pattern of transfers associated with terrorism. Likewise, the mean sensitivity (precision) is higher for grouped data, which means that the model is more precise in detecting terrorist attacks when they indeed happen. However, the mean specificity (recall) drops considerably, translating to a larger number of false alarms. However, as mentioned above, 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 terrorist attack is accurate 99.37% of the time, but the balanced accuracy for this guess is 49.68%. Using this measure of performance, we observe that the RF model has better performance at the flagged user and in the *Gambling* group but is not significantly different from a naive guess in the *Exchange* group and has worse performance in the *Service* group. Most of the anomalous transfers flagged by the model (approximately 370 transfers) correspond to the Sri Lanka Easter bombing; however, we also successfully predict several other events in the period.

Detecting terrorist financing is challenging due to the expertise of terrorist financiers in camouflaging money transfers. However, blockchain transparency may enable learning from transaction patterns to detect future similar behavior. Although a simple blockchain-based model to predict terrorist attacks has several shortcomings, the performance of these models is better than a naive guess. The performance further improves in a group of select users. These results are helpful for agencies to incorporate blockchain inputs into their anti-terrorist attack models. Additionally, the

results provide a baseline approach for market players to detect on-chain suspicious activities.

8. Conclusion

The proliferation of blockchain-based cryptocurrencies, which essentially use public accounting ledgers, may impede the significant efforts of governments, market regulators, and financial institutions to curtail illicit activities. However, the public blockchain ledger ensures transparency of the flows of funds previously observed only by the involved parties. We demonstrate, using accounting knowledge and forensic accounting techniques, that the behavior of users and transfers can be exploited to detect and predict large-scale, highly visible terrorist attacks.

However, the predictive ability of on-chain transfers is subject to several limitations. First, the predictive accuracy may improve if the training set contains known transfers associated with terrorist events—for instance, in the case where the researcher has off-chain information about specific terrorist wallets. Second, the results are limited to a group of large Bitcoin service providers. Model accuracy likely increases with the number of mapped wallets used in the learning and validation sets. Third, our model predicts large-scale events one day before they happen. Agencies working to foil large-scale terrorist attacks may need more time to implement anti-terrorist operations. Finally, our results are limited to Bitcoin-like ecosystems.

While the evolution of blockchain analytics, the increasing scrutiny by regulators, and the rise of more opaque cryptocurrencies (such as Monero and ZCash) suggest that the usage of Bitcoin for illicit purposes is likely to decrease in the future, the increasing accessibility of Bitcoin around the world suggests otherwise. In addition, the launch of decentralized exchange platforms may raise the hurdle to prevent illicit cross-border transfers and money laundering. Therefore, it will not be surprising if the need for improved prediction models of on-chain transfers for illicit activities, in general, and for terrorist financing, in particular, increases in the near future.

REFERENCES

- ABARBANELL, J. S., and B. J. BUSHEE. "Fundamental Analysis, Future Earnings, and Stock Prices." *Journal of Accounting Research* 35 (1997): 1–24.
- ACHARYA, A. *Targeting Terrorist Financing: International Cooperation and New Regimes* (Contemporary Terrorism Studies), First edition. Oxfordshire: Routledge, 2012.
- ALOOSH, A., and J. LI. "Direct Evidence of Bitcoin Wash Trading." Working paper, Neoma Business School and George Mason University 2021. Available at <https://doi.org/10.2139/ssrn.3362153>.
- ALTMAN, E. I. "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy." *Journal of Finance* 23 (1968): 589–609.
- AMIRAM, D.; E. LYANDRES; and D. RABETTI. "Competition and Product Quality: Fake Trading in Crypto Exchanges." Working paper, Tel Aviv University 2021. Available at <https://doi.org/10.2139/ssrn.3745617>.

- ASHCROFT, J., and J. W. SNOW. National Money Laundering Strategy. U.S. Treasury, 2003. Available at <https://home.treasury.gov/system/files/136/archive-documents/js10102js1010.pdf>.
- BAMBER, L. S. "The Information Content of Annual Earnings Releases: A Trading Volume Approach." *Journal of Accounting Research* 24 (1986): 40–56.
- BAMBER, L. S. "Unexpected Earnings, Firm Size, and Trading Volume Around Quarterly Earnings Announcements." *The Accounting Review* 62 (1987): 510–32.
- BEAVER, W. H. "Market Prices, Financial Ratios, and the Prediction of Failure." *Journal of Accounting Research* 6 (1968): 179–92.
- BELASCO, A.; M. EAGLEN; L. HARTIG; T. JONAS; M. MCCORD; and J. MUELLER. "Counterterrorism Spending: Protecting America While Promoting Efficiencies and Accountability." *Conventional Defense, Budgeting for Foreign Affairs and Defense* (2018). Available at https://www.stimson.org/wp-content/files/file-attachments/CT_Spending_Report_0.pdf.
- BERNARD, V. L. "Cross-Sectional Dependence and Problems in Inference in Market-based Accounting Research." *Journal of Accounting Research* 25 (1987): 1–48.
- BERNARD, V. L., and J. K. THOMAS. "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?" *Journal of Accounting Research* 27 (1989): 1–36.
- BORGERS, M. J. "Regulating and Combating Underground Banking." *Criminal Law Forum* 20 (2009): 97–111.
- BOURVEAU, T.; E. T. DE GEORGE; A. ELLAHIE; and D. MACCIOCCHI. "The Role of Disclosure and Information Intermediaries in an Unregulated Capital Market: Evidence from Initial Coin Offerings." *Journal of Accounting Research* 60 (2021): 129–67.
- BROWN, S. J., and J. B. WARNER. "Using Daily Stock Returns: The Case of Event Studies." *Journal of Financial Economics* 14 (1985): 3–31.
- BUENCAMINO, L., and S. GORBUNOV. "Informal Money Transfer Systems: Opportunities and Challenges for Development Finance." United Nations, DESA Discussion Paper No. 26, 2002.
- BUSHEE, B. J. "The Influence of Institutional Investors on Myopic R&D Investment Behavior." *The Accounting Review* 73 (1998): 305–33.
- CAMPBELL, C. J., and C. E. WESLEY. "Measuring Security Price Performance Using Daily NASDAQ Returns." *Journal of Financial Economics* 33 (1993): 73–92.
- CASCINO, S.; M. CORREIA; and A. TAMAYO. "Does Consumer Protection Enhance Disclosure Credibility in Reward Crowdfunding?" *Journal of Accounting Research* 57 (2019): 1247–302.
- CHENG, S. F.; G. DE FRANCO; H. JIANG; and P. LIN. "Riding the Blockchain Mania: Public Firms' Speculative 8-K Disclosures." *Management Science* 65 (2019): 5901–13.
- CONG, L. W.; W. R. LANDSMAN; E. L. MAYDEW; and D. RABETTI. "Tax-Loss Harvesting with Cryptocurrencies." Working paper, Cornell University, UNC-Chapel Hill, and Tel Aviv University 2022. Available at <https://doi.org/10.2139/ssrn.4033617>.
- CONG, L. W.; X. LI; K. TANG; and Y. YANG. "Crypto Wash Trading." Working paper, 2021. Available at <https://doi.org/10.2139/ssrn.3530220>.
- COOK, D. M., and T. SMITH. "The Battle for Money Transfers: The Allure of PayPal and Western Union over Familial Remittance Networks." *Journal of Information Warfare* 10 (2011): 19–36.
- COPELAND, T. E. "Liquidity Changes Following Stock Splits." *Journal of Finance* 34 (1979): 115–41.
- CORE, J. E.; W. R. GUAY; S. A. RICHARDSON; and R. S. VERDI. "Stock Market Anomalies: What Can We Learn from Repurchases and Insider Trading?" *Review of Accounting Studies* 11 (2006): 49–70.
- CORMEN, T. H.; C. E. LEISERSON; R. L. RIVEST; and C. STEIN. *Introduction to Algorithms*, Second edition. Cambridge: MIT Press, 2001.
- CORRADO, C. "The Financial Reporting Environment: Review of the Recent Literature." *Journal of Accounting and Economics* 50 (1989): 385–95.
- CORRADO, C. J., and T. L. ZIVNEY. "The Specification and Power of the Sign Test in Event Study Hypothesis Tests Using Daily Stock Returns." *The Journal of Financial and Quantitative Analysis* 27 (1992): 465–78.

- CUCULIZA, C.; C. ANTONIOU; A. KUMAR; and A. MALIGKRIS. "Terrorist Attacks, Analyst Sentiment, and Earnings Forecasts." *Management Science* 64 (2021): 2579–608.
- DAI, J.; P. BYRNES; Q. LIU; and M. VASARHELYI. "Audit Analytics: A Field Study of Credit Card After-Sale Service Problem Detection at a Major Bank." *Rutgers Studies in Accounting Analytics*. edited by J. Dai, M. A. Vasarhelyi, A. F. Medinets. In: Audit analytics in the financial industry. Bingley: Emerald Publishing Limited, 2019: SBN 978-1-78743-086 17–33.
- DEYOUNG, K., and D. FARAH. Infighting Slows Hunt for Hidden al Qaeda Assets; Funds Put in Untraceable Commodities. *Washington Post*, 2002. <https://www.washingtonpost.com/archive/politics/2002/06/18/infighting-slows-hunt-for-hidden-al-qaedaassets/ded68cf6-b290-41a8-b078-8147a84afdbc/>.
- DION-SCHWARZ, C.; D. MANHEIM; and P. B. JOHNSTON. *Terrorist Use of Cryptocurrencies: Technical and Organizational Barriers and Future Threats*. Santa Monica, CA: RAND Corporation, 2019.
- FOLEY, S.; J. R. KARLSEN; and T. J. PUTNINS. "Sex, Drugs, and Bitcoin: How Much Illegal Activity is Financed Through Cryptocurrencies?" *The Review of Financial Studies* 32 (2019): 1798–853.
- FUSARO, T., and M. HOUGAN. Presentation to the US Securities and Exchange Commission. *Bitwise Asset Management*, 2019. File No. SR-NYSEArca-2019-01. <https://www.sec.gov/comments/sr-nysearca-2019-01/srnysearca201901-5164833-183434.pdf>.
- GODESELL, D.; M. WELKER; and N. ZHANG. "Earnings Management During Antidumping Investigations in Europe: Sample-wide and Cross-sectional Evidence." *Journal of Accounting Research* 55 (2017): 407–57.
- GRIFFIN, J. M., and A. SHAMS. "Is Bitcoin Really Un-Tethered?" *Journal of Finance* 74 (2020): 1913–64.
- HENG, Y., and K. McDONAGH. *Risk, Global Governance and Security: The Other War on Terror*. *Global Security Studies*, First edition. New York: Routledge, 2009.
- ICAEW. BLOCKCHAIN AND THE FUTURE OF ACCOUNTANCY. *ICAEW Thought Leadership, IT Faculty*. London: Institute of Chartered Accountants in England and Wales, 2018.
- IRWIN, A., and G. MILAD. "The Use of Crypto-Currencies in Funding Violent Jihad." *Journal of Money Laundering Control* 19 (2016): 407–25.
- JAMES, G.; D. WITTEN; T. HASTIE; and R. TIBSHIRANI. *An Introduction to Statistical Learning: with Applications in R*. *Springer Texts in Statistics*, Seventh edition. Berlin: Springer, 2017.
- KANG, C.; C. LEE; K. KO; J. WOO; and J. HONG. "De-Anonymization of the Bitcoin Network Using Address Clustering." *Blockchain and Trustworthy Systems*. BlockSys 2020. *Communications in Computer and Information Science*, 1267. Singapore: Springer, 2020: 489–501.
- KAPPOS, G.; H. YOUSAF; M. MALLER; and S. MEIKLEJOHN. "An Empirical Analysis of Anonymity in Zcash." The 27th USENIX Security Symposium, 2018. Available at <https://arxiv.org/abs/1805.03180>.
- KATISIRI, R. "Hamas Raises Bitcoin Donations via US Crypto Exchange." *Hardfork*, 2019. <https://thenextweb.com/hardfork/2019/04/26/hamas-bitcoin-donations-address/>.
- KISER, S. *Financing Terror: An Analysis and Simulation for Affecting Al Qaeda's Financial Infrastructure*. Santa Monica, CA: RAND Corporation, 2005.
- KRANACHER, M. J., and R. RILEY. *Forensic Accounting and Fraud Examination*, Second edition. Hoboken, NJ: Wiley, 2019.
- KUHN, M., and K. JOHNSON. *Applied Predictive Modeling*, Second edition. Berlin: Springer, 2018.
- LANDIS, J. R., and G. G. KOCH. "The Measurement of Observer Agreement for Categorical Data." *Journal of Accounting and Economics* 33 (1977): 159–74.
- LAPTEV, N.; S. AMIZADEH; and I. FLINT. "Generic and Scalable Framework for Automated Timeseries Anomaly Detection." *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2015): 1939–47.
- LENNART, A. "Bitcoin Transactions, Information Asymmetry and Trading Volume." *Quantitative Finance and Economics* 4 (2020): 365–81.
- LEVITT, S., and S. VENKATESH. "An Economic Analysis of a Drug-Selling Gang's Finances." *The Quarterly Journal of Economics* 115 (2000): 755–89.
- LIMODIO, N. "Terrorism Financing, Recruitment and Attacks." Forthcoming, *Econometrica* (2022). Working Paper, 2021. Available at <https://pedl.cepr.org/publications/terrorism-financing-recruitment-and-attack>.

- LYANDRES, E.; B. PALAZZO; and D. RABETTI. "ICO Success and Post-ICO Performance." *Forthcoming, Management Science* (2021). <https://doi.org/10.1287/mnsc.2022.4312>.
- MAKAROV, I., and A. SCHOAR. "Trading and Arbitrage in Cryptocurrency Markets." *Journal of Financial Economics* 135 (2020): 293–319.
- MALKIN, L., and Y. ELIZUR. "Terrorism's Money Trail." *World Policy Journal* 19 (2002): 60–70.
- MCNEMAR, Q. "Note on the Sampling Error of the Difference Between Correlated Proportions or Percentages." *Psychometrika* 12 (1947): 153–57.
- MEIKLEJOHN, S.; M. POMAROLE; G. JORDAN; K. LEVCHENKO; D. MCCOY; G. M. VOELKER; and S. SAVAGE. "A Fistful of Bitcoins: Characterizing Payments Among Men with no Names." *Communications of the ACM* 59 (2018): 86–93.
- MICHEL, J. "Do Unverifiable Disclosures Matter? Evidence from Peer-to-Peer Lending." *The Accounting Review* 87 (2012): 1385–413.
- NAKAMOTO, S. "Bitcoin: A Peer-to-Peer Electronic Cash." White paper, 2008. Available at <https://bitcoin.org/bitcoin.pdf>.
- NAVIAS, M. "Finance Warfare and International Terrorism." *Political Quarterly* 73 (2002): 57–79.
- OHLSON, J. A. "Financial Ratios and the Probabilistic Prediction of Bankruptcy." *Journal of Accounting Research* 18 (1980): 109–31.
- PIETH, M. *Financing of Terrorism: Following the Money*. European Journal of Law Reform 4. 2002: 365–376.
- PIOTROSKI, J. D. "Value Investing: The Use of Historical Financial Statement Information to Separate Winners from Losers." *Journal of Accounting Research* 38 (2000): 1–41.
- PUKELSHEIM, F. "The Three Sigma Rule." *The American Statistician* 48 (1994): 88–91.
- ROGERS, J. L.; D. J. SKINNER; and S. L. C. ZECHMAN. "The Role of the Media in Disseminating Insider-trading News." *Review of Accounting Studies* 21 (2016): 711–39.
- RON, D., and A. Shamir. "Quantitative Analysis of the Full Bitcoin Transaction Graph." *Financial Cryptography and Data Security* edited by A. R. Sadeghi. Lecture Notes in Computer Science 7859, Darmstadt, Germany: Springer, (2013): 6–24.
- RUDNER, M. "Hizbullah Terrorism Finance: Fund-Raising and Money-Laundering." *Studies in Conflict & Terrorism* 33 (2010): 700–15.
- SCHOTT, P. A. *Reference Guide to Anti-Money Laundering and Combating the Financing of Terrorism*, Second edition. Washington, DC: The World Bank, International Monetary Fund, 2006.
- SINGLETON, T. W. and A. J. Singleton. *Fraud Auditing and Forensic Accounting*, Fourth edition. New York: NY John Wiley & Sons, 2010.
- SLOAN, R. G. "Do Stock Prices Fully Reflect Information in Accruals and Cash Flows About Future Earnings?" *The Accounting Review* 71 (1996): 289–315.
- SOKOLOV, K. "Ransomware Activity and Blockchain Congestion." *Journal of Financial Economics* 141 (2021): 771–82.
- TASCA, P.; A. HAYES; and S. LIU. "The Evolution of the Bitcoin Economy: Extracting and Analyzing the Network of Payment Relationships." *The Journal of Risk Finance* 19 (2018): 94–126.
- XIA, T., and Y. GU. "Building Terrorist Knowledge Graph from Global Terrorism Database and Wikipedia." *IEEE International Conference on Intelligence and Security Informatics* (2019): 194–196. <https://doi.org/10.1109/ISI.2019.8823450>.
- YERMACK, D. "Is Bitcoin a Real Currency? An Economic Appraisal." *Handbook of Digital Currency* (Elsevier, 2015): 31–44.
- ZMIJEWSKI, M. E. "Methodological Issues Related to the Estimation of Financial Distress Prediction Models." *Journal of Accounting Research* 22 (1984): 59–82.
- ZORN, C. "A Solution to Separation in Binary Response Models." *Journal of Political Analysis* 13 (2005): 157–70.