BRL
Blockchain Research Lab

BRL Working Paper Series No. 7

Dominating OP Returns: The Impact of Omni and Veriblock on Bitcoin

Elias Strehle [1,*], Fred Steinmetz [1,2]

[1] Blockchain Research Lab, Colonnaden 72, 22303 Hamburg, Germany
[2] University of Hamburg, Faculty of Business, Economics & Social Sciences,
Von-Melle-Park 5, 20146 Hamburg, Germany

[*] Correspondence: strehle@blockchainresearchlab.org

Published: 10 Mar 2020

Abstract: Bitcoin has always been used to store arbitrary data, particularly since Bitcoin Core developers added a dedicated method for data storage in 2014: the OP Return operator. This paper provides an in-depth analysis of all OP Return transactions published on Bitcoin between September 14, 2018, and December 31, 2019. The 32.4 million OP Return transactions (22% of all Bitcoin transactions) published during this period added 10 GB to the blockchain's size. Almost all OP Return transactions can be attributed to one of 37 blockchain services. The two dominant services are Veriblock (58% of OP Return transactions) and Omni/Tether (40%). Veriblock transactions pay only 14% of the average transaction fee, partly because most of them are submitted during times when overall activity on Bitcoin is low. Omni transactions, on the other hand, pay more than twice the average transaction fee and therefore compete with regular Bitcoin transactions for inclusion in new blocks.

Keywords: Blockchain; Econometrics of Cryptocurrencies; OP_RETURN; Tether; Transaction Fees

1 Introduction

The Bitcoin blockchain was created to serve a single purpose: the operation of a decentralized and secure digital currency. But even its inventor(s) Satoshi Nakamoto could not resist using it for something else. The following text is hidden in the details of the first Bitcoin transaction: "The Times 03/Jan/2009 Chancellor on brink of second bailout for banks." Whether intended as a political statement or simply as a timestamp, the text goes beyond the declared purpose of the Bitcoin blockchain as a store of payment-related data. Many other Bitcoin users have since found it useful to publish nonpayment data on the blockchain.

Most early methods of data insertion on Bitcoin had the undesirable effect of permanently increasing the number of unspent transaction outputs (UTXO), forcing full nodes to keep a growing number of "fake" transactions in memory. Bitcoin Core[1] developers reacted in 2014 and released the OP Return script operator. It allows users to add up to 80 bytes of arbitrary data to a transaction without

[1] Bitcoin Core is the reference implementation for Bitcoin clients. Originating from Satoshi Nakamoto's (2008) first client, the source code for Bitcoin Core is open source. By 2016, more than 400 developers had contributed to the enhancement and continued upgrading of the Bitcoin system (Antonopoulos, 2016).

negatively affecting the UTXO. However, the release notes emphasize that the introduction of the OP Return operator "is not an endorsement of storing data on the blockchain. Storing arbitrary data in the blockchain is still a bad idea."[2]

But is it? Bitcoin's perceived immutability, censorship-resistance and accessibility make it an excellent medium for publishing data in a transparent and tamper-proof way. By mid-2017, more than two million OP Return transactions had been published on Bitcoin, mostly by organizations in search of a sustainable business model (Bartoletti et al., 2019). The overall impact however was small at the time. OP Return transactions accounted for only one percent of all transactions.

This changed in the following years. From the end of 2017 until the end of 2019, the total number of OP Return transactions grew from four million to thirty-nine million.[3] Almost all new OP Return transactions came from two services: Omni and Veriblock. Omni is the protocol layer that supports the well-known "stable coin" Tether, whose value is pegged to the US Dollar. Veriblock is a blockchain project built around a novel security mechanism called Proof-of-Proof (PoP), which combines notarization of the current state of Veriblock's blockchain with a complex system of mining rewards. Veriblock emerged from obscurity in 2019 as it became the single largest source of OP Return transactions, accounting for up to 20% of daily Bitcoin transactions.[4]

Both Omni and Veriblock sparked controversy for storing large amounts of "arbitrary" data on Bitcoin. Advocates lauded them for using Bitcoin in an innovative way and raising the demand for its token.[5] Opponents suspected them of raising transaction fees, crowding out regular Bitcoin transactions, and illegitimately bloating the size of the blockchain.[6]

While various papers have analysed nonpayment data on Bitcoin, none of them covers the recent surge in the number of OP Return transactions. The most comprehensive studies (Matzutt et al., 2018; Bartoletti et al., 2019) only cover OP Return transactions published until 2017 and thus account for less than ten percent of all OP Return transactions published by the end of 2019. In particular, we are not aware of any scientific studies that discuss the sudden rise of Veriblock and its effect on Bitcoin. This paper aims to fill the gap with an in-depth analysis of all OP Return transactions published on Bitcoin between September 14, 2018 (the day of the first Veriblock transaction on Bitcoin), and December 31, 2019. We identify which services published OP Return transactions, how much they paid for them and how they contributed to the size of the Bitcoin blockchain. Our main focus lies on Omni and Veriblock as the two dominant publishers of OP Return transactions on Bitcoin.

2 Storing data on Bitcoin

2.1 Methods of data insertion, the UTXO set, and the OP Return operator

The desirable characteristic of Bitcoin's blockchain is the authenticity of its records (Lemieux, 2017). The cumulative computational effort of creating and maintaining the Bitcoin blockchain presents a virtually insurmountable security wall. Incentivization schemes ensure that the users work hard to protect the integrity of the data submitted to the blockchain, which can therefore be considered immutable. This property has induced users to find ways to insert arbitrary data ever since the beginning of the Bitcoin blockchain. Arbitrary data on Bitcoin refers to any data or file that is not related to

[2] https://bitcoin.org/en/release/v0.9.0#opreturn-and-data-in-the-block-chain.

[3] https://opreturn.org/op-return-per-month.

[4] https://www.forbes.com/sites/ktorpey/2019/01/09/a-new-blockchain-project-is-generating-20-of-daily-bitcoin-transactions.

[5] https://www.veriblock.org/#qanda, Question "How does VeriBlock help Bitcoin?"

[6] See e.g. the sources cited in https://news.bitcoin.com/veriblock-captured-close-to-60-of-btcs-op-return-transactions-in-2019.

either payment or the operation of the Bitcoin network.

As data can only be added to the blockchain in the course of transactions, any arbitrary data must be submitted as part of the transaction data. Bitcoin's scripting language *Script* defines standard script types, which principally frame the playing field for creating data insertion techniques.[7] Sward et al. (2018) differentiate such methods according to whether they rely on Bitcoin's standard script type Pay-to-Script-Hash (P2SH). Methods that do not use P2SH include Coinbase transactions, Pay-to-Fake-Key (P2FK), Pay-to-Fake-Key-Hash (P2FKH), OP Return and Pay-to-Fake-Multi-Signature (P2FMS). Methods that use P2SH are more complex; they include Pay-to-Fake-Script-Hash (P2FSH), Data Drop and Data Hash. In the most widely used methods, various parameters within input and output scripts are substituted or modified to convey specific content in the form of files (including images) or plain text.

Matzutt et al. (2018) distinguish methods for data insertion according to whether they utilize input or output scripts. Input scripts are part of every Bitcoin transaction (except Coinbase transactions) and provide a "proof" of satisfying a set of conditions outlined in a previous output. Consequently, spending this output in a subsequent transaction requires the input script to satisfy the conditions of the previous output script. Output scripts relate to signing transactions, so that an input script allows a miner to verify that the sender is also the owner of the respective BTC (Matzutt et al., 2016).

As an example, the P2FKH method uses the P2PKH standard script to insert arbitrary data in a special field for the hash value of the public key that contains the destination address of a transaction. This field is part of a transaction's output script and requires a certain amount of BTC to be attached. By replacing an existing public key with a hash of arbitrary data, the BTC associated with the transaction's output become unspendable, which adds to the cost of using this method. However, the hash that replaces the key is permanently stored by miners because they do not know whether the hash corresponds to an existing public key (Sward et al., 2018). Accordingly, submitting arbitrary data via Bitcoin transactions always requires getting miners to verify and include those transactions in blocks. Because the miners might refuse to support transactions that aim to insert arbitrary data, methods have been developed to disguise that purpose.

The insertion of arbitrary transaction data may require miners to store more data. Besides inflating Bitcoin's blockchain, several methods of inserting arbitrary data affect another important data set: the Unspent Transaction Outputs (UTXO) set. The UTXO set "is a subset of Bitcoin transaction outputs that have not been spent at a given moment" (Delgado-Segura et al., 2019). It is used for validating new transactions without the necessity to inspect the whole blockchain. Miners and full nodes in the Bitcoin network must continuously update the UTXO set. Every transaction produces changes in the UTXO set as transaction inputs consume UTXOs and outputs generate new ones. According to Sward et al. (2018), five out of nine methods for inserting arbitrary data affect the UTXO set and thus produce additional overhead for miners.[8] Bartoletti and Pompianu (2017) find that until 2017, approximately 1.5% of the transactions that were initiated for inserting data into Bitcoin's blockchain bloat the UTXO set.

The OP Return operator presents an alternative to those problematic data insertion techniques. It was introduced with the release of Bitcoin's reference client Bitcoin Core 0.9. Although the description of

[7] A miner can also decide to accept non-standard transactions which include arbitrary data. However, Bitcoin's reference client only accepts standard transactions based on the standard scripts. Matzutt et al. (2016) identify approximately 290 000 non-standard transactions in Bitcoin's blockchain until July 2016, almost all of which are OP Return transactions with an empty payload. Only 132 are non-standard transactions that do not use standard script templates. We therefore consider the relevance of non-standard transactions negligible.

[8] For a comprehensive analysis of the different techniques and their implications for Bitcoin, see Sward et al. (2018), Cremona et al. (2019), and Bartoletti et al. (2019).

the version release specifies that the introduction of the OP Return operator is not an endorsement for storing data on Bitcoin's blockchain, its recognition in the reference client is most likely a reaction to the increasing activity of storing arbitrary data on the blockchain using various inefficient methods Antonopoulos (2016).[9]

The OP Return is an instruction in Bitcoin's scripting language. It can be used in the locking script of an output. In a transaction, one output at most may contain the OP Return instruction. The interpreter of the scripting language ignores any data that follows it. Nevertheless, the arbitrary data is part of the transaction and will eventually be stored in the blockchain. The data cannot exceed 80 bytes[10], which is much less than with other methods of data insertion (Sward et al., 2018).

The instruction to ignore the data differentiates the OP Return from other "more harmful" techniques. It produces a prunable output, which means that Bitcoin nodes which store the complete blockchain—a requirement for full nodes and mining nodes—can decide whether to store the output of the OP Return transaction in their UTXO set. This accelerated its acceptance by Bitcoin miners. Furthermore, unlike other methods, OP Return transactions do not require the sender to burn BTC valued higher than "dust"[11] (Matzutt et al., 2016; Sward et al., 2018). That is, the amount of Bitcoin associated with the output of the OP Return transaction can be small or even zero. This is especially attractive for protocols whose services require storing arbitrary data on Bitcoin because it reduces costs compared to other methods. It also reduces complexity as transactions can be initiated solely for the purpose of inserting data via OP Returns. In sum, Sward et al. (2018) conclude that the OP Return is the most efficient method for inserting small amounts of data.

While the OP Return operator has been part of Bitcoin's scripting language since the beginning (Bartoletti and Pompianu, 2017), it was only its recognition in Bitcoin's reference client in March 2014 that paved the way for its acceptance and increased usage. Since then, the OP Return has experienced widespread adoption, having already become the most widely used data insertion method by mid-2015 (Matzutt et al., 2016).

2.2 Existing research on OP Return use and service categorization

2.2.1 Number and share of OP Return transactions

The share of OP Return transactions relative to other data insertion methods increased significantly since its introduction in 2014, which led researchers to investigate it. Existing research focuses either on data insertion in Bitcoin in general or on OP Return transactions specifically. The investigated data periods range from Bitcoin's first ever mined block in 2009 until November 2018, with transaction counts of up to 356.6 million, depending on the study period and methodological differences.

Matzutt et al. (2016) classify the content submitted to Bitcoin's blockchain. The transaction data they analyse spans the period from Bitcoin's beginning in 2009 until July 2016. Of the 146 million Bitcoin transactions they examine, approximately 936 000 are OP Return transactions.[12] For the period until February 2017, Bartoletti and Pompianu (2017) identify approximately 1.9 million OP Return transactions. Thus, the number of such transactions doubled between mid-2016 and early 2017. While Matzutt et al. (2016) find that only 0.64% of all transactions include OP Return instructions, the share

[9] https://www.coindesk.com/bitcoin-core-dev-update-5-transaction-fees-embedded-data.

[10] This limit has changed over time. For a historic review of the OP Return, see Bartoletti and Pompianu (2017).

[11] According to the Bitcoin Core reference implementation, "dust" is the output of a transaction whose redemption fee exceeds one third of its value (see https://github.com/bitcoin/bitcoin/blob/v0.10.0rc3/src/primitives/transaction.h#L139-L146).

[12] The authors refer to OP Return transactions as "nulldata" transactions. These terms are used synonymously, see e.g. https://bitcoin.org/en/glossary/null-data-transaction.

Table 1: OP Return transactions in Bitcoin's blockchain as identified in the literature.

Paper	**Data period**		**Transactions (Mio.)**	
	Date	**Blocks**	**OP Return**	**Total**
Matzutt et al. (2016)	01.09.09 – 31.07.16	0 – 423 075	0.9	146.0
Bartoletti and Pompianu (2017)	01.09.09 – 15.02.17	0 – 453 200	1.9	196.6
Bartoletti et al. (2019)	01.09.09 – 10.08.17	0 – 480 000	2.9	–
Matzutt et al. (2018)	01.09.09 – 31.08.17	0 – 482 869	3.1	250.9
Bistarelli et al. (2019)	01.09.09 – 14.11.18	0 – 550 000	> 5.8	356.6
Faisal et al. (2018)	29.03.13 – 06.07.17	228 596 – 474 451	3.7	–
This paper	14.09.18 – 31.12.19	541 306 – 610 681	32.4	147.6

increased to 0.96% by February 2017 (Bartoletti and Pompianu, 2017). The latter authors report a share of OP Return transactions from their start in March 2014 until February 2017 of 1.16%.

Matzutt et al. (2018) analyse the impact of arbitrary content on Bitcoin based on transaction data from July 2009 until August 2016. The authors identify 3.5 million transactions that contain nonpayment data, 86.8% of which are OP Return transactions. With a reported share of 1.2% of all Bitcoin transactions, the authors' result for the share OP Return transactions significantly exceeds those of the aforementioned works. Faisal et al. (2018) obtain even higher values. The authors analyse a dataset containing all blocks of the Bitcoin blockchain from the end of March 2013 to the beginning of July 2017 and identify a total of 3.7 million OP Return transactions. Compared to Bartoletti and Pompianu (2017), the number of OP Return transactions appears to have almost doubled within less than five months in 2017.

Bartoletti et al. (2019) identify approximately 2.9 million OP Return transactions in 480 000 blocks from 2009 until August 2017. This amounts to about 1.2% of all Bitcoin transactions since 2009 and about 1.4% of all transactions since March 2014. In light of Bartoletti and Pompianu (2017), these results imply that approximately one million OP Return transactions were initiated within less than six months from February 15, 2017, until August 10, 2017. To the best of our knowledge, the latest publication on OP Return transactions is Bistarelli et al. (2019), which analyses standard and non-standard transactions in Bitcoin until November 2018. For the last year, the authors report that the number of OP Return transactions exceeded five million, which would imply another significant increase since mid-2017.

These differences either indicate rapid growth in the number of OP Return transactions throughout 2017 and 2018 or result from different approaches to data collection and the identification of OP Return transactions. Table 1 sums up these results from the literature.

2.2.2 Cumulative data load, OP Return sizes and costs

Given the relevance of OP Return transactions for inserting arbitrary data in Bitcoin (Matzutt et al., 2018), past research has investigated the actual data load submitted through OP Return transactions. According to Bartoletti and Pompianu (2017), the total size of all arbitrary data inserted into the Bitcoin blockchain via OP Return transactions amounted to 44 MB (1.9 million transactions) by early 2017, while Bitcoin's blockchain itself amounted to 102 GB at the time. Overall, those OP Return transactions consumed 323 MB. Matzutt et al. (2018) find that OP Return and Coinbase transactions jointly account for 81% (96 MB) of the nonpayment data on Bitcoin's blockchain. For the 2.9 million

OP Return transactions published by mid-2017, Bartoletti et al. (2019) calculate a total data load of 77 MB. Roughly 10% of all OP Return transactions had no data attached.

Bistarelli et al. (2019) analyse the size of data inserted through OP Return transactions. They find that most transactions (four million) have a size of only 20 bytes, while half a million are 40 bytes and one million are 80 bytes in size. This distribution of OP Return sizes is mostly due to the changes in the OP Return capacity over the course of subsequent releases of the Bitcoin Core reference client. Approximately 297 000 OP Return transactions did not carry any data load, 222 000 of which can be attributed to a network stress test during September 2015.

Interestingly, only BTC 2 615 were burned while submitting data to Bitcoin's blockchain in the time period considered by Bistarelli et al. (2019). With the introduction of the OP Return operator in Bitcoin's reference client, burning Bitcoin is no longer a prerequisite for inserting data. Therefore, the more interesting indicator is the amount of fees paid for OP Return transactions. Sward et al. (2018) report that the total cost of storing 80 bytes of arbitrary data on Bitcoin's blockchain using OP Return transactions amounts to 6 340 satoshi/byte (one BTC equals 100 million satoshi), although it is unclear which data period was considered to obtain this number.

We assume that most transactions that carry arbitrary data are initiated solely for the purpose of inserting this data into the Bitcoin blockchain, and thus do not carry any meaningful transactions of BTC between users. Accordingly, when measuring the data load of OP Return transactions in Bitcoin, one should consider the size of the entire transaction instead of the size of the inserted data.

2.2.3 Service providers and protocol categories

OP Return transactions can often be associated with specific protocols by identifiers contained in the OP Return data. Most protocol owners and inventors are companies that use OP Return transactions to offer specific services. The identifiers are structures in the data of OP Returns that allow users and service providers to find and verify data submitted to the blockchain. Based on these recurring structures different types of services have been identified. In their analysis of approximately 1.9 million OP Return transactions, Bartoletti and Pompianu (2017) identify 22 protocols which account for 51% of the transactions.[13] The authors categorize these protocols as assets, document notary, digital arts, and others. Matzutt et al. (2018) categorize such services as digital notary services, secure releases of cryptographic commitments, and non-equivocation schemes. Faisal et al. (2018) in turn categorize these services as key value stores, notary/doc, assets, any message, and proof of ownership. Bartoletti et al. (2019) find 45 protocols and 39 identifiers.[14] The authors categorize these protocols as financial, notary, digital rights management, message, and subchain. The definition of the financial category is roughly identical to definitions of the asset category in other research articles. In sum, the different taxonomies overlap considerably. As the last categorization of protocols is the finest, we will expand upon it in more detail.

The purpose of financial protocols such as Colu (Colu Technologies DLT Ltd., 2018) and Omni (Willett et al., 2017) is the management of assets other than Bitcoin. Associated OP Return transactions include information about the ownership and size of the transferred asset. This way, a currency is implemented through a technical layer "on top" of Bitcoin, without the need for a dedicated blockchain. Notary protocols provide proof of existence and timestamps for documents by submitting cryptographic identifiers to the blockchain. This way they attest that a document existed at the time of submission to the blockchain, and the unchanged integrity of a file since its submission can be verified. In combination with private key encryption, the originator of the file can also be ascertained:

[13] The authors identify 25 protocols, but only provide data for 22.

[14] Some protocols use more than one identifier and 19 protocols do not use any identifiers.

it is the owner of the private key at the time of submission. Examples of notary protocols are Stampery (Sánchez de Pedro Crespo and García, 2017) and Factom (Snow et al., 2014). Digital rights management protocols facilitate the management of access rights to documents and files. As an example of such services, Monegraph focuses on enabling creators of digital art to maintain control over the reproduction and use of their work through technical licensing methods. Similar services include Binded (formerly Blockai) and Ascribe. Message services like Eternity Wall and BitAlias store text messages and aliases on Bitcoin's blockchain. Finally, Subchain protocols "construct transaction chains to record execution traces of third-party smart contracts" (Bartoletti et al., 2019). The only service in this category is the Blockstack Naming Service (BNS, formerly Blockstore).

The categories indicate that OP Return transactions are used for various purposes and across different domains. According to Bartoletti et al. (2019), "financial" is the dominant category in terms of the number of transactions and data load. From about 2 000 in April 2015, the number of transactions in the financial category increased to approximately 200 000 in July 2017, accounting for 63% of all OP Return transactions. Financial protocols (referred to as asset protocols in other research articles), also appear to be the most common category in previous articles, accounting for 27% (Bartoletti and Pompianu, 2017) and 30% (Faisal et al., 2018) of the transactions. Matzutt et al. (2016) find that 10% of all non-empty OP Return transactions are associated with the protocol Open Assets, which is also an asset protocol.

To the best of our knowledge, previous comprehensive studies of OP Return protocols only incorporate blockchain data until August 2017. Since then, the landscape of protocols and their use have changed drastically. During the surge in cryptocurrency prices in late 2017 and early 2018, the protocol Omni gained significant importance. Moreover, new protocols have been developed which secure the blockchains of alternative cryptocurrencies. Among these, Veriblock received the most media attention due to the large number of transactions associated with it.[15] It appears that no published research covers the current state of OP Return transactions, the shares of the different protocols and the allegedly dominant services, Omni and Veriblock.

2.2.4 The dominant services: Omni and Veriblock

Omni is a protocol layer built on top of Bitcoin. It was developed to enable the creation of new tokens on the Bitcoin blockchain (Willett et al., 2017). The most popular token that uses the Omni protocol is the "stable coin" Tether (USDT). Its originator Tether Ltd. promises that all USDT in circulation are backed by a reserve and can be redeemed at any time, which should guarantee that one USDT is always worth one USD (Tether Ltd., 2016).

Every Bitcoin transaction can carry an Omni transaction. Conversely, every Omni transaction requires a Bitcoin transaction. The Omni protocol itself does not charge transaction fees for basic transactions, but the sender must pay the transaction fees for the corresponding Bitcoin transaction. A higher transaction fee typically implies quicker publication of the corresponding transaction.

Figure 1a shows the OP Return format of a typical Omni transaction. It begins with the identifier "omni" and includes the transaction type (e.g. a "simple send"), the targeted property (e.g. Tether) and the transfer amount. The OP Return data does not include any addresses because Omni reuses the sender and receiver addresses of the Bitcoin transaction itself.

The Omni protocol has been in operation since 2013. Tether was launched on Omni and thus on Bitcoin in 2014 and remains the main use case for Omni, although Tether is by now also available on other blockchains, including Ethereum. Still, Omni has been one of the most important services

[15] See e.g. https://news.bitcoin.com/veriblock-captured-close-to-60-of-btcs-op-return-transactions-in-2019.

$$\texttt{OP\_RETURN}\ \underbrace{\texttt{6F 6D 6E 69}}_{\text{Omni prefix}}\ \underbrace{\texttt{00 00}}_{\text{Version}}\ \underbrace{\texttt{00 00}}_{\text{Type}}\ \underbrace{\texttt{00 00 00 1F}}_{\text{Property ID}}\ \underbrace{\texttt{00 00 00 A2 8A 73 F8 D3}}_{\text{Amount}}$$

(a) An Omni transaction. The Omni prefix translates to "omni" in ASCII. The transaction is a "simple send" ("Type 0") in Tether ("Property ID 31") for an amount of USDT 6 981.075 580 99. For further details see https://jochen-hoenicke.de/crypto/omni.

$$\texttt{OP\_RETURN}\ \underbrace{\texttt{00 0A 5C BC}}_{\text{Block height}}\ \underbrace{\texttt{00 02}}_{\text{Version}}\ \underbrace{\texttt{20 30 66 F9 58 DB DD AA 85 6B EC 12}}_{\text{Previous block hash}}$$

$$\underbrace{\texttt{F8 36 FA F4 7E D1 66 3E 5B}}_{\text{First prev. keystone hash}}\ \underbrace{\texttt{DD 9F 60 9F 6C A7 6E 88 6E}}_{\text{Second prev. keystone hash}}$$

$$\underbrace{\texttt{63 EF 5D 5E 4B 24 E2 57 57 EC 31 1F BA 91 C9 8C}}_{\text{Merkle root hash}}\ \underbrace{\texttt{5D E2 ED 52}}_{\text{Timestamp}}$$

$$\underbrace{\texttt{06 4A 2E DD}}_{\text{Difficulty}}\ \underbrace{\texttt{29 A3 EF 1D}}_{\text{Nonce}}\ \underbrace{\texttt{01 34 E1 B9 97 8F C5 04 57 33 C4 80 A3 D9 92 E6}}_{\text{Miner address}}$$

(b) A Veriblock transaction. The hashes and the miner address are shortened to fit the 80 byte limit of Bitcoin's OP Return code. The endorsement references the Veriblock block with height 679 100 on the mainnet (the testnet is "Version 1", the mainnet is "Version 2") and was timestamped to 2019-11-30, 23:29:38 UTC. For further details see Section 8.6 of VeriBlock, Inc. (2019).

Figure 1: OP Return formats of Omni and Veriblock transactions.

using Bitcoin's OP Return transactions: more than 40% of all OP Return transactions ever published on Bitcoin are Omni transactions.[16]

Veriblock is a blockchain project that is built around a novel security mechanism called Proof-of-Proof (PoP).[17] Veriblock's users can act as "Proof-of-Proof miners" by publishing a snapshot of the Veriblock blockchain in an OP Return transaction on Bitcoin. In the context of Proof-of-Proof, the published snapshot is referred to as an endorsement. While new blocks on Veriblock are created by standard Proof-of-Work (PoW) mining, forks are resolved in favour of the chain that has more and older endorsements on Bitcoin. The intention is to make Veriblock just as secure against 51% attacks as Bitcoin, in that a successful attack against Veriblock would require control over the endorsements published on Bitcoin. Consequently, Veriblock advertises that cryptocurrencies built on top of its blockchain inherit Bitcoin's security without having to operate on Bitcoin directly.

Publishing endorsements on Bitcoin entails transaction fees. To create an incentive to pay these fees, Veriblock grants a mining reward to every user who publishes an endorsement and references it in a special transaction on Veriblock. In a sense, Proof-of-Proof miners exchange BTC (in the form of transaction fees) for Veriblock's cryptocurrency VBK (in the form of mining rewards). The first users to endorse a new block on Bitcoin receive the largest share of the mining reward. Proof-of-Proof miners thus face the challenge of finding the optimal level of transaction fees that guarantees speedy publication on Bitcoin at minimal cost and thus the best "exchange rate" between BTC and VBK.

Figure 1b shows the OP Return format of a Veriblock transaction. It was designed to fit the 80 byte limit and contains the height of the endorsed Veriblock block, the protocol version, a timestamp, and

[16] https://opreturn.org/op-return-protocols; the category "6f6d6e" corresponds to Omni transactions.

[17] Information on Veriblock was obtained from the original whitepaper (Sanchez and Fisher, 2018), the most recent version of the whitepaper (VeriBlock, Inc., 2019), and Veriblock's website (https://veriblock.org) and wiki (https://wiki.veriblock.org). The authors clarified some details via chat in the Veriblock Discord channel under the username "estrehle".

multiple shortened hashes, including the hash of the most recent block on Veriblock and hashes of the two most recent “keystone blocks.”[18] The endorsement also contains the Veriblock address of the miner who published it.

Veriblock’s Proof-of-Proof mining on the Bitcoin mainnet began on September 14, 2018, with the launch of the Veriblock “high-noon” testnet. The testnet ran until March 3, 2019, and was succeeded by the Veriblock mainnet on March 25, 2019. Outstanding balances on the testnet were doubled and transferred to the mainnet. The first services began testing integration with Veriblock in early 2020.[19] Accordingly, Veriblock’s own blockchain saw little activity until the end of 2019, averaging approximately one thousand non-mining transactions per day between September 2018 and November 2019.[20] This is in sharp contrast to the activity of Veriblock’s Proof-of-Proof miners on Bitcoin, who were estimated to publish more than 77 000 transactions per day in 2019.[21]

3 Actual use of OP Return transactions

3.1 Methodology

To study the current use of OP Return transactions and the impact of Omni and Veriblock, we analyse all transactions published on the Bitcoin mainnet between September 14, 2018 (the day of the first Veriblock transaction on Bitcoin), and December 31, 2019. We downloaded all blocks mined during this period (block heights 541 306 to 610 681) via the Smartbit Bitcoin Block API, ignoring orphaned blocks. We classify a transactions as an OP Return transaction if exactly one of its outputs is of the “nulldata” type, i.e. an output whose locking script begins with the OP Return operator. Note that Coinbase transactions can be OP Return transactions.

To categorize the OP Return transactions, we check its data for identifiers. An identifier is a character sequence included in the OP Return data to allow its attribution to a protocol. Most identifiers are prefixes to the OP Return data. For example, OP Return data published in accordance with the Omni protocol begins with “omni”. The Komodo protocol is an exception: it uses a suffix instead of a prefix. We check all identifiers from Table 2 in Bartoletti et al. (2019) and add identifiers that we discovered by visual inspection of a subset of our data. Four of these additional identifiers could not be attributed to a known protocol but are included nonetheless: “DC-L5”, “POR”, “POTX”, and “VX”. One identifier, the “SegWit commitment”, is only used by Bitcoin miners in Coinbase transactions as part of the segregated witness consensus layer. In total, we check for 50 identifiers associated with 37 services.

Veriblock’s OP Return data does not contain an identifier, so we use a heuristic to detect Veriblock transactions: If an OP Return transaction could not be matched with a known identifier and has exactly 80 bytes of OP Return data, we parse the data according to the format described in Figure 1b and extract the values for version, height and timestamp. The transaction is categorized as a Veriblock transaction if either of the following sets of conditions applies:

[18] Keystone blocks are intended as protection against hidden long-range attacks on Proof-of-Proof. For details, see Section 6.2 of VeriBlock, Inc. (2019).

[19] Two of these services are Placeholder (https://www.placeh.io/index.jsp#map) and Pexa (https://medium.com/@ryanmhein/pexa-project-committed-to-incorporating-the-veriblock-proof-of-proof-consensus-protocol-27f073aab4de/).

[20] Data obtained from Veriblock’s Block API (https://explore.veriblock.org/api/block/{block_height}). The API documentation can be found at https://wiki.veriblock.org/index.php/Dashboard_API.

[21] https://www.forbes.com/sites/ktorpey/2019/01/09/a-new-blockchain-project-is-generating-20-of-daily-bitcoin-transactions.

(i) *version* equals 1 (Veriblock testnet), *height* between 0 and 472 000 (the highest block mined on the testnet), and *timestamp* between 1 536 883 200 (September 14, 2018; our start date) and 1 553 472 000 (March 25, 2019; the day the mainnet was launched);

(ii) *version* equals 2 (Veriblock mainnet), *height* between 0 and 761 487 (the highest block mined during our data period), and *timestamp* between 1 553 472 000 (see above) and 1 577 836 799 (December 31, 2019; our end date).

Since these conditions are not sufficient, there is a possibility of overestimating the number of Veriblock transactions. However, given the specificity of the conditions, we expect the number of misclassified transactions to be negligible. An exact categorization would require elaborate cross-validation of OP Return data with the Veriblock blockchain.

All remaining OP Return transactions are categorized as "Empty" if the OP Return data comprises 0 bytes and as "Unknown" otherwise.

For each Bitcoin block, we obtain its height and timestamp. We further calculate aggregate metrics for the following groups of transactions: all transactions in the block, all OP Return transactions, and each category of OP Return transactions (i.e. all protocols as distinguished by identifiers, plus Veriblock, Empty, and Unknown). The aggregate metrics are: number of transactions, total transaction fee (in satoshi), size (in bytes), and virtual size (in vbytes). Size refers to the total size of a transaction, not the size of its OP Return data. In addition, we calculate the transaction fee per virtual size (in satoshi/vbyte) as the most common measure of the value of a transaction to Bitcoin miners.

We additionally obtain the following information for every Omni and Veriblock transaction: Block height and timestamp, time when the transaction was first seen by the Smartbit Bitcoin server (as a biased estimate of when the transaction was submitted to the Bitcoin network), number of inputs and outputs, input amount, output amount sent to non-input addresses, fee, size, virtual size, and the OP Return data.

3.2 Data analysis

Between September 14, 2018, and December 31, 2019, 147.6 million transactions were published on the Bitcoin mainnet, adding 73 GB to blockchain size and paying BTC 22 000 in transaction fees—roughly USD 150 million (at 7 000 USD/BTC). Our sample comprises 32.4 million OP Return transactions (22% of all transactions), with a total size of 10 GB (14%) and total transaction fees of BTC 3 300 (15%). Note that we measure the total size of OP Return transactions, not the size of OP Return data, which explains the discrepancy with previous studies.

We are able to categorize almost 99% of the OP Return transactions based on an identifier or on our heuristic for Veriblock transactions. Fewer than 391 000 transactions are "Unknown". Table 2 at the very end of the paper contains the results of our general analysis of OP Return protocols, including a comparison between our data period and the period considered by Bartoletti et al. (2019), which ends in August 2017. Our data confirm the dominance of Veriblock (58% of all OP Return transactions) and Omni (40%). Only five protocols (Veriblock, Omni, Factom, Komodo, and Blockstore) averaged more than 1 000 transactions per month—compared to nine services in the period considered by Bartoletti et al. (2019). Ten protocols seem to have ceased operation altogether. Among them is Counterparty, formerly the largest publisher of OP Return transactions.

One key metric is the transaction fee per virtual size. Since the Bitcoin protocol limits block size in terms of virtual size, miners generally maximize their income by including those transactions in their candidate blocks that carry the highest fee per virtual byte. The average transaction in our data period paid 37.98 satoshi/vbyte. The average OP Return transaction paid 33.91 satoshi/vbyte, but fees

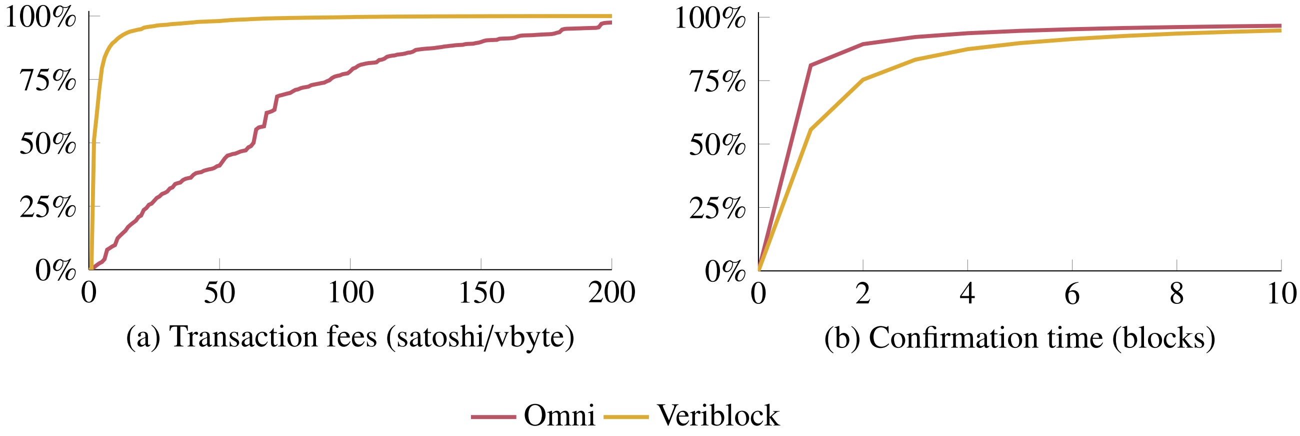

Figure 2: Distribution of transaction fees and confirmation time for Veriblock and Omni transactions.

vary strongly between protocols. The unknown protocols “POR” and “POTX” paid particularly high average fees of more than 184 satoshi/vbyte, while Factom paid only 9.96 satoshi/vbyte. How the service RSK managed to publish more than 6 700 transactions without paying any transaction fees at all remains an open question.

The most surprising observations regarding transaction fees concern Veriblock and Omni. Veriblock transactions paid an average fee of 5.45 satoshi/vbyte—only about 14% of the overall average. Consequently, Veriblock transactions account for only 1% of all transaction fees paid during our data period, despite making up 13% of the transactions. By contrast, Omni transactions paid an average fee of 67.92 satoshi/vbyte, which is almost twice the average. Figure 2a plots the distribution of Veriblock’s and Omni’s transaction fees. The fees for Omni transactions are also more dispersed than the fees for Veriblock. Almost all Veriblock transactions paid less than 20 satoshi/vbyte, while 24% of Omni transactions paid less than 20 satoshi/vbyte and 21% paid more than 100 satoshi/vbyte.

We are not aware of previous attempts to explain the unusually low fees paid by Veriblock or the unusually high fees paid by Omni. In the following we consider two potential partial explanations: Omni users pay high fees to ensure quick confirmation of their transactions; Veriblock users pay low fees because they are most active when overall activity on Bitcoin is low.

In principle, we define the confirmation time of a transaction as the number of blocks published between submission to the Bitcoin network and inclusion in a block. If a transaction is included in the first block after submission, the confirmation time is one block; if it is included in the second block after submission, the confirmation time is two blocks; etc. However, we do not know the exact submission time of transactions—that information is generally only available to the publisher of the transaction. Therefore, we replace the time of submission with a biased proxy, namely the time when the transaction was first seen by our data source, the Smartbit Bitcoin server. Figure 2b plots the distribution of confirmation times and shows that Omni transactions are indeed confirmed more quickly than Veriblock transactions. For Omni, 81% of the transactions are confirmed in the first block and 92% are confirmed by the third block, compared to 56% and 83% for Veriblock. In this regard at least, the higher transaction fees paid by Omni users pay off.

The Bitcoin network is most active on working days and during European daytime. High activity typically implies high transaction fees, as more transactions compete for confirmation. Figure 3 shows the anticyclical behaviour of Veriblock’s publishers. They are most active on the weekend and during European night-time, which allows them to publish transactions at low cost. However,

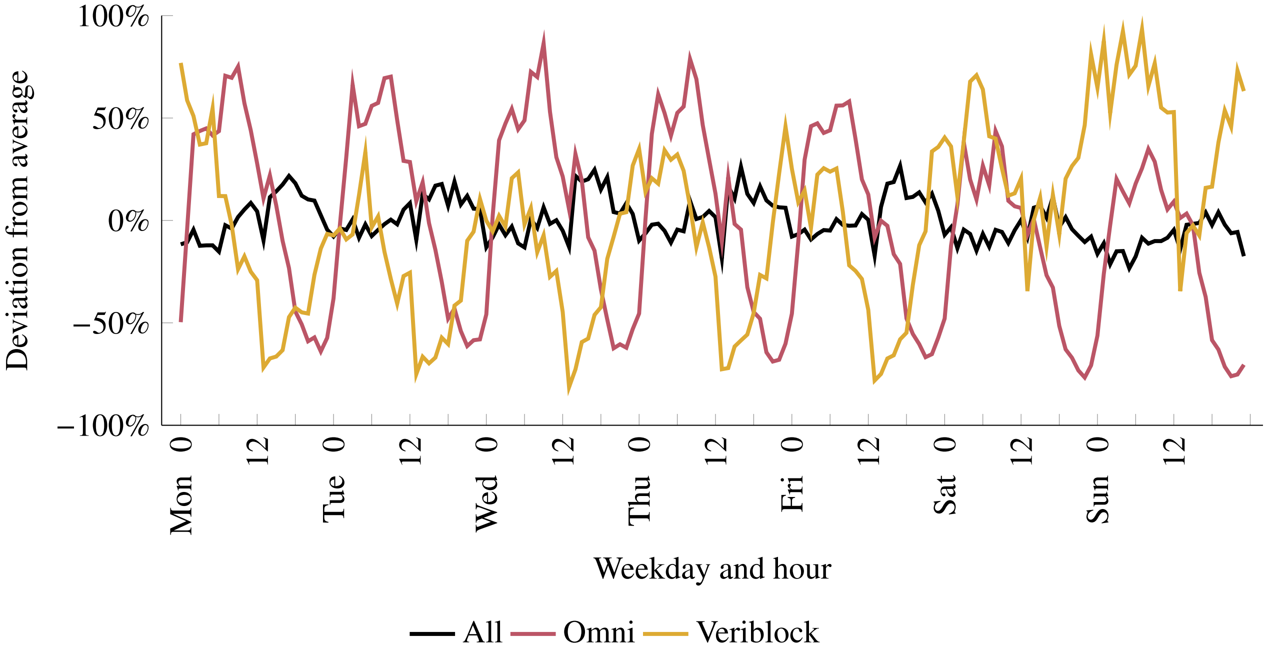

Figure 3: Number of transactions by weekday and hour (UTC). Veriblock transactions are mostly published on the weekend and during European night-time, when overall activity on Bitcoin is low. Omni transactions are mostly published on working days during European night-time and early daytime.

publishers of Omni transactions, too, tend to avoid the times of greatest activity and publish most transactions during European night-time.

For further insights on who publishes Omni and Veriblock transactions, we parse the OP Return data of all Omni and Veriblock transactions according to the formats shown in Figures 1a and 1b. As described in Section 2.2.4, Veriblock transactions are published by Proof-of-Proof miners who compete for mining rewards. We find that the 18.9 million Veriblock transactions on Bitcoin mention 905 different Veriblock addresses, so this is the maximum number of miners or mining pools that participated in Proof-of-Proof mining. Most of them ceased mining before mid-2019. Veriblock transactions published between July and December 2019 mention only 77 Veriblock addresses. During this time, the 30 most active addresses account for 93% of all Veriblock transactions.

Our analysis suggests that the vast majority of Veriblock transactions served no other purpose than the publication of OP Return data: Approximately 92% of the Veriblock transactions are self-transfers that send a single unspent transaction output to the same address, subtracting only the transaction fee and adding the Veriblock OP Return output.

It is worth noting that the 18.9 million Veriblock transactions on Bitcoin contain only 5.9 million different endorsements; on average, a Proof-of-Proof miner publishes the same endorsement on Bitcoin more than three times. This is a consequence of Veriblock's reward scheme, which pays a higher share to miners who publish the same endorsement multiple times.

More than 700 tokens[22] use the Omni protocol, but 98% of the Omni transactions in our data period are simple send transactions for Tether (USDT). Most transactions transfer between USDT 100 and USDT 10 000. Surprisingly, the fees of a transaction are independent of its transfer size. Even transactions below USDT 100 carry the unusually high fees observed for all Omni transactions.

[22] https://omniexplorer.info/properties/production.

While the share of self-transfers is lower for Omni than for Veriblock, 17% of the Omni transactions do not send any BTC to a non-input address, i.e. an address that is not contained in any of the transaction inputs, and 94% only send 10 000 satoshi (approximately USD 0.70) or less. This suggests that most Omni transactions, too, serve no purpose beyond data publication.

3.3 Discussion

Almost six years have passed since Bitcoin Core developers released the OP Return operator as a dedicated method of publishing nonpayment data on the blockchain. Our analysis confirms that OP Return transactions have become an integral part of the Bitcoin ecosystem. In fact, the comparison with past studies reveals that the number of OP Return transactions increased more than five-fold from November 2018 until the end of 2019. Discarding the publication of nonpayment data on Bitcoin as a "bad idea" does not do this reality justice.

The comparison with Bartoletti et al. (2019) reveals major shifts in the field of OP Return protocols. Many protocols have disappeared since 2017, others have become more popular. The fact that Omni and Veriblock make up 98% of all OP Return transactions in our data period should not conceal the success of other protocols like Factom, Komodo, and Blockstore. The dominance of Omni and Veriblock in terms of the sheer number of transactions is mostly a result of their respective designs: Every Omni transaction requires a Bitcoin transaction, and Veriblock pays a reward for every Proof-of-Proof endorsement that is published on Bitcoin. It would be premature to conclude that Omni and Veriblock are "more successful" than other OP Return protocols—other protocols are responsible for far fewer Bitcoin transactions but might carry just as much business value. Komodo, for instance, makes very similar security promises to Veriblock but appears to implement them with a much smaller number of Bitcoin transactions.[23] Similarly, a single transaction published by a notary service might secure assets worth thousands of dollars.

It is safe to conclude, however, that Omni and Veriblock are the only OP Return protocols with a significant effect on the Bitcoin ecosystem in terms of blockchain size and transaction fees. Omni transactions alone account for more than 13% of the total transaction fees paid during our data period. While miner income is currently still dominated by mining rewards, large and reliable sources of transaction fees will become increasingly important for Bitcoin miners as mining rewards continue to halve every four years.

Protocols that pay high transaction fees are advantageous to miners, but not necessarily to other Bitcoin users. Both Omni and Veriblock have been accused of raising transaction fees and increasing confirmation times for "real" Bitcoin transactions (as opposed to OP Return transactions). We are unable to issue a final verdict, but it seems likely that the large number and high fees of Omni transactions have increased the transaction fees that regular Bitcoin transactions have had to pay to ensure swift confirmation. Omni's OP Return transactions can thus impact Bitcoin's intended functionality as a payment infrastructure.

Veriblock is an altogether different matter. Very low transaction fees, relatively long publication times and the anticyclical behaviour of its Proof-of-Proof miners suggest that Veriblock transactions are more aptly regarded as a supplement to the regular use of Bitcoin: they fill up the free space in Bitcoin blocks that is left over by other transactions. More than 90% of the Veriblock transactions pay a transaction fee of 10 satoshi/vbyte or less and are thus no serious competition for most other transactions. As long as this remains the case, Veriblock transactions are unlikely to have a significant effect on the level of Bitcoin transaction fees during times of high activity. In the future, it is possible that Veriblock's mining mechanism becomes unprofitable as overall activity on Bitcoin increases and

[23] https://github.com/SuperNETorg/komodo/wiki/Delayed-Proof-of-Work-(dPoW)-Whitepaper.

the publication of transactions at minimal cost becomes impossible. Yet we may also see the value of Veriblock's cryptocurrency increase significantly, which could entice Proof-of-Proof miners to pay much higher transaction fees on Bitcoin.

The assumption that most OP Return transactions are initiated solely for the purpose of data publication is confirmed for Omni and Veriblock. We therefore suggest that any analysis of the effects of OP Return protocols on Bitcoin consider the size of the entire transaction, not just the size of the published data. With a joint volume of 10 GB over our data period, transactions associated with the Omni and Veriblock protocols certainly have an impact on Bitcoin's blockchain size. The question remains to what extent these OP Return transactions displaced regular Bitcoin transactions.

Unlike Proof-of-Work mining, profitable Proof-of-Proof mining does not require any specialized hardware. Any user of Veriblock and Bitcoin can compete for mining rewards. Veriblock itself encourages Bitcoin users to "mine on the side" by adding an OP Return output in the appropriate format to their regular transactions.[24] We find no evidence of this happening during our data period; almost all Veriblock transactions exist for the sole purpose of earning mining rewards. Veriblock's Proof-of-Proof mining is conducted by a small and shrinking number of dedicated Proof-of-Proof miners.

4 Conclusion and further research

The total number of OP Return transactions on Bitcoin has increased almost sixfold during the period from September 14, 2018 to December 31, 2019. We find that 22% of the transactions from our data period are OP Return transactions, and 98% of these are associated with either Omni or Veriblock. While Omni, which has Tether as its most important token, has been around since 2014 and has achieved a significant market share in relation to other protocols, Veriblock is a novel protocol whose properties incentivize its users, so-called Proof-of-Proof miners, to actively publish OP Return transactions on Bitcoin. Accounting for 58% of all OP Return transactions, Veriblock is the dominant service in terms of the number of transactions and the associated data load.

Compared to earlier studies, the landscape of protocols has changed drastically. We find that many protocols that published considerable numbers of OP Return transactions in the past, such as Counterparty (Bartoletti and Pompianu, 2017) and Open Assets (Matzutt et al., 2016), have ceased to operate. In particular, we note the reduced activity of notary services.

Our analysis shows that OP Return transactions have a considerable impact on Bitcoin's intended functionality as a payment structure. Most OP Return transactions are initiated solely for the purpose of inserting arbitrary data. Their cumulative data load, considering the size of the entire transaction and not just the arbitrary data, exceeds 10 GB during the period we investigate. Yet the impacts of Veriblock and Omni differ. While Veriblock transactions are anticyclical to network activity and thus to high fees in Bitcoin, Omni transactions are not. This suggests that Omni transactions typically compete with regular Bitcoin transactions, while Veriblock transactions typically do not. Omni transactions also have an impact on transaction fees, while Veriblock transactions rather fill blocks during times of low network activity. Omni transactions carry fees of 67.92 satoshi/vbyte, which is almost twice the average fee. By contrast, Veriblock transactions only carry fees of 5.45 satoshi/vbyte.

It is unclear why Omni users pay such high transaction fees, especially for low-value transfers of Tether. The official Omni wallet contains an algorithm that attempts to compute optimal transaction fees.[25] However, given that USDT are often used to buy other cryptocurrencies, it is likely that the

[24] https://www.veriblock.org/#qanda, Question "How does VeriBlock help Bitcoin wallet providers?"

[25] https://github.com/OmniLayer/omniwallet/wiki/How-to-adjust-the-miner-fee.

majority of Tether and thus Omni transactions are sent from and to exchanges.[26] Users of exchanges might be less knowledgeable about appropriate transaction fees on Bitcoin and more focussed on short confirmation times. Exchanges might act accordingly by charging high fees to their users and then sending overpriced transactions to the Bitcoin network.

Veriblock's Proof-of-Proof mechanism is a novel approach for decentralizing notarization services. After initial popularity, the number of Proof-of-Proof miners has decreased. Between July and December 2019, only 77 Veriblock addresses were referenced in Veriblock transactions, with the 30 most active ones accounting for 93% of all Veriblock transactions. Consequently, most mining rewards on Veriblock are paid out to very few individuals. It remains to be seen whether this centralization of funds will negatively affect the credibility and economics of Veriblock. Additionally, the anticyclical activity of Proof-of-Proof mining indicates that its level of protection varies inversely with activity on the Bitcoin mainnet. In times of high Bitcoin activity, Veriblock's incentive mechanism may open windows for specific attacks on the Veriblock blockchain and the cryptocurrencies which use it. How this interrelates with the price volatility of VBK and BTC and variations in average transaction fees on Bitcoin is a topic for future research.

On the other hand, the anticyclical behaviour currently exhibited by Proof-of-Proof miners on Bitcoin may catch on. Not all Bitcoin transactions require swift publication. For example, notarization services or exchanges that rebalance their accounts may be flexible enough to postpone transactions by hours or even days in order to achieve significantly lower transaction fees, just as Proof-of-Proof miners manage to reduce their transaction fees to only 14% of the average value. It certainly seems worthwhile to further explore the effect of transaction timing on transaction fees.

As long as Bitcoin is perceived as a secure public data store, services will use it to store non-payment data. Our analysis reveals that different services affect Bitcoin in different ways. Omni users pay unusually high transaction fees, which increases the income of Bitcoin miners but increases competition for regular Bitcoin transactions. Veriblock users withdraw from competition by paying low transaction fees and publishing transactions when activity on Bitcoin is low and blocks contain empty space. Metaphorically speaking, Veriblock transactions fill Bitcoin blocks in the same way that padding fills a parcel. Note however that the "social cost" of increasing the storage requirements for every Bitcoin full node is just as pronounced for Veriblock as it is for Omni. It remains to be seen how Omni and Veriblock fare as demand for Bitcoin evolves. Nonetheless, we conclude that OP Return transactions have become a major factor in the Bitcoin ecosystem and are likely to remain so for the foreseeable future.

References

Antonopoulos, A. M. (2016). *Mastering Bitcoin: Unlocking Digital Cryptocurrencies*. O'Reilly Media, Inc., second edition. ISBN: 978-1-4919-5437-9.

Bartoletti, M., Bellomy, B., and Pompianu, L. (2019). A journey into Bitcoin metadata. *Journal of Grid Computing*, 17(1):3–22. https://doi.org/10.1007/s10723-019-09473-3.

Bartoletti, M. and Pompianu, L. (2017). An analysis of Bitcoin OP_RETURN metadata. In Brenner, M., Rohloff, K., Bonneau, J., Miller, A., Ryan, P. Y. A., Teague, V., Bracciali, A., Sala, M., Pintore, F., and Jakobsson, M. (eds.), *Financial Cryptography and Data Security*, pp. 218–230, Cham. Springer International Publishing. https://doi.org/10.1007/978-3-319-70278-0_14.

Bistarelli, S., Mercanti, I., and Santini, F. (2019). An analysis of non-standard transactions. *Frontiers in Blockchain*, 2:7. https://doi.org/10.3389/fbloc.2019.00007.

[26] As shown for example by the trading volumes of BTC/USDT compared to BTC/USD on https://www.bti.live/bitcoin-coin.

Colu Technologies DLT Ltd. (2018). *Colu local network: Version 1.0* [Whitepaper]. Retrieved February 27, 2020, from https://cln.network/pdf/cln_whitepaper.pdf.

Cremona, K., Tabone, D., and De Raffaele, C. (2019). Cybersecurity and the blockchain: Preventing the insertion of child pornography images. In O'Conner, L. (ed.), *2019 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC)*, Guilin, China. IEEE. https://doi.org/10.1109/CyberC.2019.00042.

Delgado-Segura, S., Pérez-Solà, C., Navarro-Arribas, G., and Herrera-Joancomartí, J. (2019). Analysis of the Bitcoin UTXO set. In Zohar, A., Eyal, I., Teague, V., Clark, J., Bracciali, A., Pintore, F., and Sala, M. (eds.), *Financial Cryptography and Data Security: FC 2018*, volume 10958 of *Lecture Notes in Computer Science*, Berlin/Heidelberg, Germany. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-58820-8_6.

Faisal, T., Courtois, N., and Serguieva, A. (2018). The evolution of embedding metadata in blockchain transactions. In *2018 International Joint Conference on Neural Networks (IJCNN)*, Rio de Janeiro, Brazil. IEEE. https://doi.org/10.1109/IJCNN.2018.8489377.

Lemieux, V. L. (2017). Blockchain and distributed ledgers as trusted recordkeeping systems: An archival theoretic evaluation framework. In *Proceedings of the 2017 Future Technologies Conference (FTC)*, pp. 41–48, Vancouver, Canada. The Science and Information (SAI) Organization.

Matzutt, R., Hiller, J., Henze, M., Ziegeldorf, J. H., Müllmann, D., Hohlfeld, O., and Wehrle, K. (2018). A quantitative analysis of the impact of arbitrary blockchain content on Bitcoin. In Meiklejohn, S. and Sako, K. (eds.), *Financial Cryptography and Data Security*, pp. 420–438, Berlin/Heidelberg, Germany. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-58387-6_23.

Matzutt, R., Hohlfeld, O., Henze, M., Rawiel, R., Ziegeldorf, J. H., and Wehrle, K. (2016). Poster: I don't want that content! On the risks of exploiting Bitcoin's blockchain as a content store. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1769–1771, New York, USA. Association for Computing Machinery. https://doi.org/10.1145/2976749.2989059.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system* [Whitepaper]. Retrieved February 27, 2020, from https://bitcoin.org/bitcoin.pdf.

Sanchez, M. and Fisher, J. (2018). *Proof-of-Proof: A decentralized, trustless, transparent, and scalable means of inheriting Proof-of-Work security: Version 1.0p* [Whitepaper]. Retrieved February 27, 2020, from https://www.veriblock.org/wp-content/uploads/2018/03/PoP-White-Paper.pdf.

Sánchez de Pedro Crespo, A. and García, L. (2017). *Blockchain timestamping architecture version 6 (BTAv6)* [Whitepaper]. Retrieved February 27, 2020, from https://doi.org/10.13140/RG.2.2.17223.80805.

Snow, P., Deery, B., Lu, J., Johnston, D., and Kirby, P. (2014). *Factom: Business processes secured by immutable audit trails on the blockchain* [Whitepaper]. Retrieved February 27, 2020, from https://github.com/FactomProject/FactomDocs/blob/master/Factom_Whitepaper.pdf.

Sward, A., Vecna, I., and Stonedahl, F. (2018). Data insertion in Bitcoin's blockchain. *Ledger*, 3. https://doi.org/10.5195/ledger.2018.101.

Tether Ltd. (2016). *Tether: Fiat currencies on the Bitcoin blockchain* [Whitepaper]. Retrieved February 27, 2020, from https://tether.to/wp-content/uploads/2016/06/TetherWhitePaper.pdf.

VeriBlock, Inc. (2019). *Proof-of-Proof and VeriBlock blockchain protocol consensus algorithm and economic incentivization specifications: Version 1.0* [Whitepaper]. Retrieved February 27, 2020, from https://www.veriblock.org/wp-content/uploads/2019/06/Proof-of-Proof_and_VeriBlock_Blockchain_Protocol_Consensus_Algorithm_and_Economic_Incentivization_v1.0.pdf.

Willett, J. R., Hidskes, M., Johnston, D., Gross, R., Schneider, M., Zathras, Dexx, and Gilligan, S. (2017). *Omni protocol specification*. Retrieved February 27, 2020, from https://github.com/OmniLayer/spec.

Table 2: OP Return protocols, identifiers, transaction counts, fees in BTC, sizes in MB and comparison of transactions per month between the data period considered by Bartoletti et al. (2019) and our data period.

Protocol	**Identifiers**[a]	**Trans-actions**	**Fee (BTC)**	**Size (MB)**	**Fee/vsize (sat./vbyte)**	**Trans./month (Bartoletti et al., 2019)**[b]	**Trans./ month**[c]
Ascribe	`ASCRIBE`	49	0.00	0.04	13.79	1 505	3
BitAlias	`BALI`	0	–	–	–	0	0
BitProof	`BITPROOF`	0	–	–	–	26	0
Blockai/Binded	`0x1F00`	3	0.00	0.00	18.78	21	0
Blocksign	`BS`	5	0.00	0.00	21.44	40	0
Blockstore	`0x5808`, `0x5888`, `0x6964`	25 808	2.67	10.33	25.84	6 444	1 637
ChainX	`ChainX:`	9 512	0.77	3.36	25.16	–	603
CoinSpark	`SPK`	5	0.00	0.00	36.13	743	0
Colu	`CC`	5 951	0.82	2.92	28.17	9 597	377
Counterparty	`CNTRPRTY`	0	–	–	–	16 563	0
CryptoCopyright	`CryptoProof-`, `CryptoTests-`	0	–	–	–	1	0
Eternity Wall	`EW`	355	0.07	0.08	89.19	–	23
Factom	`FA`, `Fa`, `FACTOM00`, `Factom!!`	67 280	1.63	16.40	9.96	2 591	4 267
Helperbit	`HB`	1	0.00	0.00	28.21	1	0
Komodo	Suffix: `KMD 0x00`	65 714	20.50	103.86	19.74	–	4 168
LaPreuve	`LaPreuve`	0	–	–	–	2	0
Monegraph	`MG`	329	0.03	0.09	36.10	2 605	21
Nicosia	`UNicDC`	0	–	–	–	1	0
Notary	`Notary`	1	0.00	0.00	8.85	5	0
Omni	`omni`	12 882 936	2 928.82	4 357.70	67.92	12 771	817 100

OpenAssets	OA	2 670	0.85	0.88	96.73	5 196	169
Openchain	OC	5	0.00	0.00	24.10	125	0
OriginalMy	ORIGMY	0	–	–	–	5	0
Photector	PEIRMOBILE.COM, PHOTECTOR.COM	1 943	0.08	0.53	21.16	–	123
po.et	POET 0x00000012	10 244	0.81	2.74	42.39	–	650
Proof of Existence	DOCPROOF	1 110	0.12	0.27	45.53	136	70
ProveBit	ProveBit	0	–	–	–	2	0
Remembr	RMBd, RMBe	0	–	–	–	1	0
RSK	RSKBLOCK:	6 749	0.00	1.98	0.00	–	428
SmartBit	SB.D	0	–	–	–	406	0
Stampd	STAMPD##	2	0.00	0.00	96.15	18	0
Stampery	S1, S2, S3, S4, S5, S6	1 561	0.10	0.42	35.95	2 536	99
SegWit commitment	0xAA21A9ED	57 303	0.00	15.94	0.00	–	3 634
VeriBlock	No identifier	18 910 466	273.32	5 361.78	5.45	–	1 199 395
Unknown protocol	DC-L5:	2 255	0.20	0.57	35.56	–	143
Unknown protocol	POR:	382	0.19	0.10	184.18	–	24
Unknown protocol	POTX:	474	0.24	0.13	184.58	–	30
Unknown protocol	VX 0x00000005	106	0.00	0.03	10.81	–	7
Empty		77	0.01	0.03	26.05	7 171	5
Unknown		390 579	27.29	126.57	22.48	15 593	24 772
Total		32 443 875	3 258.53	10 006.74	33.91	84 103	2 057 751

[a] Characters prefixed with 0x are in hexadecimal. Spaces in identifiers should be ignored. For example, the hexadecimal identifier of po.et is 0x504F455400000012.

[b] Based on Tables 2 and 5 in Bartoletti et al. (2019). Defined as 30 * Transactions / Number of days between first item and August 10, 2017.

[c] Defined as 30 * Transactions / Number of days between September 14, 2018 and December 31, 2019.

Declarations

Availability of data and materials

The datasets used during in this paper are available from the corresponding author on request.

Conflicts of interest

Not applicable.

Funding

Not applicable.

Acknowledgments

The authors thank Sönke Häseler for proofreading the manuscript.

About the Blockchain Research Lab

The Blockchain Research Lab promotes independent science and research on blockchain technologies and the publication of the results in the form of scientific papers and contributions to conferences and other media. The BRL is a non-profit organization aiming, on the one hand, to further the general understanding of the blockchain technology and, on the other hand, to analyze the resulting challenges and opportunities as well as their socio-economic consequences.

www.blockchainresearchlab.org