

Working Paper presented at the

Peer-to-Peer Financial Systems

2020 Workshop

2020

Ride the Lightning: Turning Bitcoin into Money

Peter Zimmerman

Federal Reserve Bank of
Cleveland

Powered by



Ride the Lightning: Turning Bitcoin into Money

Abstract

We show that recent technological innovations have improved the efficiency of Bitcoin as a means of payment. We find a robust and significant association between reduced blockchain congestion since the beginning of the 2018, and adoption of the Lightning Network, a means of netting payments off the blockchain. This improvement cannot be explained by other factors, such as changes in speculative demand for Bitcoin. Our findings have implications for the design of central bank digital currencies. We show that the Lightning Network has become increasingly centralised, with payments channelled through relatively few intermediaries. We conclude that improved functioning of Bitcoin is positive for welfare, and may reduce the environmental footprint of Bitcoin mining.

JEL classification: D4, E42, G10, O33.

Keywords: Bitcoin, blockchain, cryptocurrency, Lightning Network, SegWit, networks, money.

1 Introduction

The intended purpose of Bitcoin is to serve as a means of payment outside of the control of centralised monetary authorities, and to maintain privacy for users (Nakamoto 2008). Since its introduction in 2008, it has grown immensely in value, but still sees relatively little use as a means of payment (Thakor 2019). One important reason is that blockchain technology imposes capacity constraints on handling transactions. Bitcoin can handle an average of only seven transactions per second across the entire system. If too many transactions need to be made, they must queue and wait for settlement. This limit is low compared to centralised payment systems such as Visa or Mastercard, which handle thousands of transactions each second. The constraint is a technological consequence of having a truly decentralised payments system that is secure against attack (Abadi and Brunnermeier 2018).

In recent years, Bitcoin developers have proposed various solutions to this so-called *scalability problem*, so that the cryptocurrency can achieve its potential as a universal payments system. Since the beginning of 2018, congestion in Bitcoin has fallen markedly. The number of transactions in the *mempool* — a list of payments waiting to be confirmed — has declined (see green line in Figure 1). To date, the average daily mempool count in 2019 is over 75% lower than in 2017. This reduction does not appear to have been driven by a fall in demand for Bitcoin transactions. Although demand did decline following the collapse of the cryptocurrency market at the beginning of 2018, the number of confirmed transactions has since picked up (see red line in Figure 1). In fact, at the time of writing, the Bitcoin blockchain is handling around 300,000 transactions per day, close to its all-time peak.¹

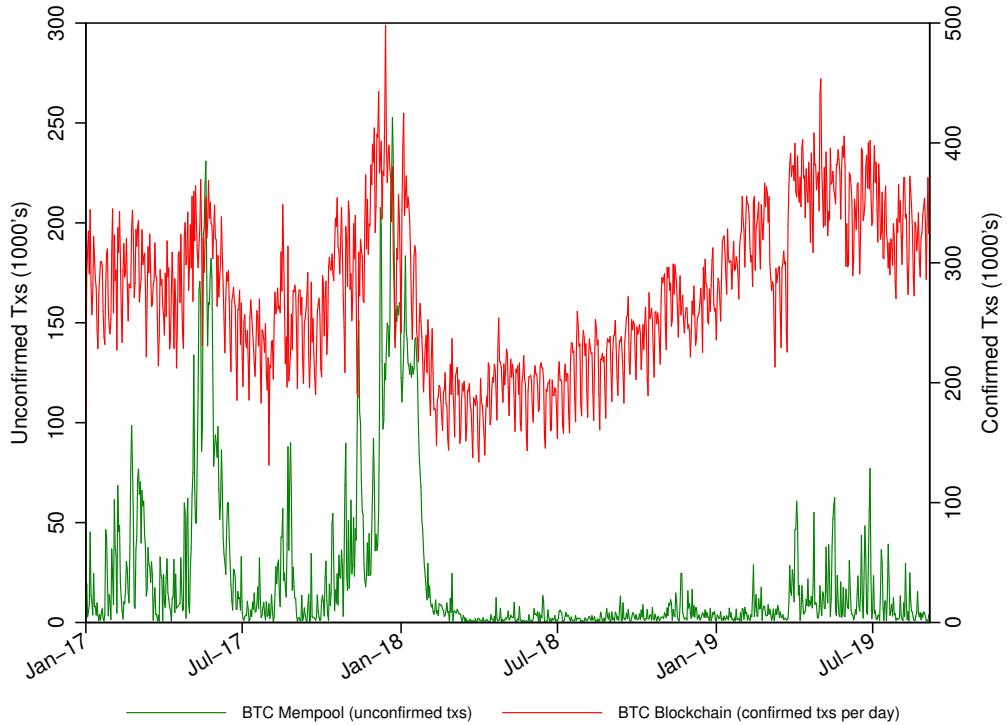
There is evidence, too, that the economic cost of congestion has fallen. Bitcoin users pay fees to miners in order to achieve settlement priority for their transactions. Figure 2 suggests that these fees have fallen since the beginning of 2018. The proportion of transactions with a fee of less than 3 satoshis per virtual byte fell from 20.1% on January 1, 2018 to 47.8% on September 5, 2019.²

¹See <https://www.blockchain.com/charts/n-transactions?timespan=all>.

²There are 100 million satoshis to a bitcoin. In this paper, “Bitcoin” refers to the entire system, while “bitcoin” (with a small ‘b’) refers to the currency unit. This is standard usage.

Figure 1: Decline in Bitcoin mempool congestion.

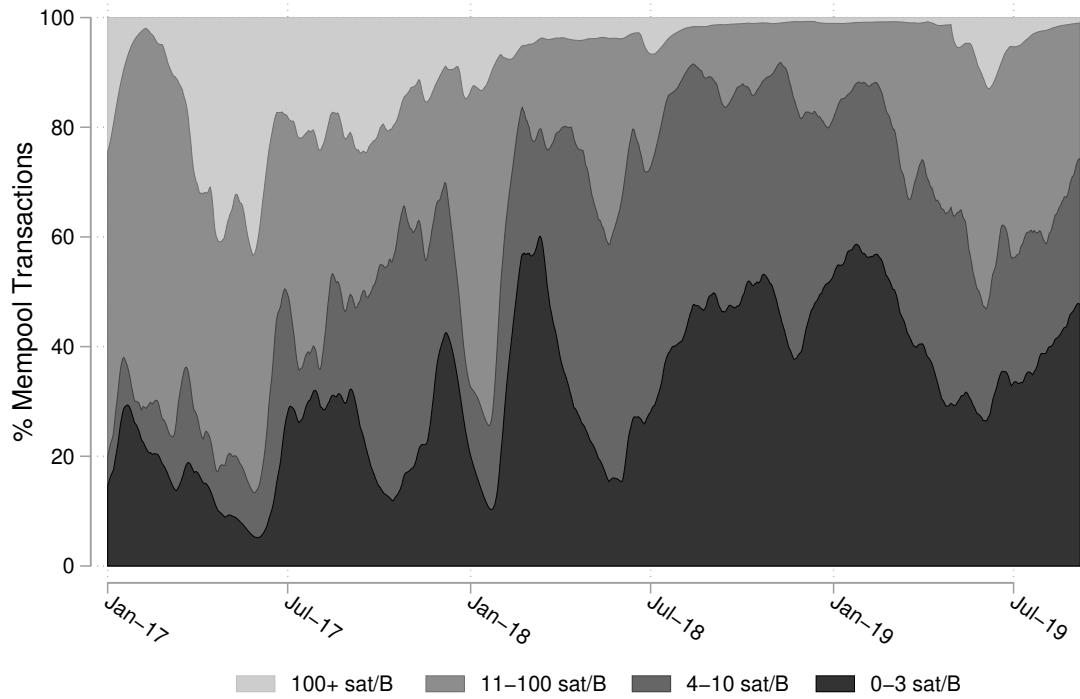
Daily data from January 1, 2017 to September 5, 2019. The red line shows the number of payments settled on the blockchain each day, while the green line shows the number waiting to be settled. Source: <https://jochen-hoenicke.de/queue>.



We show the decline in Bitcoin congestion is chiefly driven by technological innovations that have increased supply, rather than by demand factors. We examine three major new innovations introduced between late 2017 and early 2018 that have improved Bitcoin’s settlement capacity. These are: *Bitcoin Cash*, a new cryptocurrency; *SegWit*, an improvement to the efficiency of how blockchain space is used; and the *Lightning Network*, a way of settling payments off-chain. Of the three, the Lightning Network has the most significant impact, both statistically and economically. This suggests that Bitcoin can achieve even greater scalability, if adoption of the Lightning Network continues. This may have positive implications for welfare. First, the Lightning Network provides Bitcoin users with an option that may allow faster settlement of transactions at lower cost. Second, the Lightning Network reduces the total size of the blockchain,

Figure 2: Distribution of fees in the Bitcoin mempool.

Daily data from January 1, 2017 to September 5, 2019. The chart plots fees in satoshis per virtual byte. There are 100 million satoshis to a bitcoin. Source: <https://jochen-hoenicke.de/queue>.



lowering computational requirements for nodes and miners. This reduces the cost of operating a node, making the network more secure. Third, lower aggregate fees reduce the incentive to devote computational power to Bitcoin mining, with a commensurate environmental benefit.

We document the evolution of the structure of the Lightning Network and explain why it has moved to a more centralised structure. The Lightning Network allows two Bitcoin users to open a direct bilateral channel, through which they make payments to one another. When the channel is closed, only a single payment for the net amount needs to be submitted to the mempool. This requires less blockchain space — and therefore lower fees — than submitting payments directly to the mempool. However, Lightning channels must be collateralised with Bitcoin, in order to protect users against counterparty default. The total amount of collateral required can be reduced by steering

payments through an intermediary. We show that the Lightning Network has become increasingly centralised, as payments are steered through a small number of highly connected intermediaries. But competitive forces should prevent the network from becoming totally centralised.

The remainder of this paper is organised as follows. Section 2 surveys the relevant literature. Section 3 describes the data that we use, and Section 4 discusses adoption of the three new technological innovations. In Section 5, we show that, of these innovations, only the Lightning Network has had a significant impact on mempool congestion. Section 6 explores how the shape of the Lightning Network has evolved, while Section 7 discusses welfare implications. Finally, Section 8 concludes.

2 Literature

This paper is related to a wide literature on the role of cryptocurrencies as monetary assets. Schilling and Uhlig (2019) model how cryptocurrency can be adopted over time, while Bolt and van Oordt (2019) consider a speculative market. Athey et al. (2016) take a reduced-form approach and study a dynamic model of adoption where the technology could fail at any time. Biais et al. (2018) find evidence that the price of bitcoin responds to news about its ability to serve as money, suggesting that usage is not purely speculative. None of these papers focus on the blockchain settlement constraint.

This paper also relates to a newly developing literature on the fee-based market for blockchain space. Huberman, Leshno, and Moalleni (2017) uses queueing theory to assess the effect of blockchain congestion on fees and waiting times. Easley, O’Hara, and Basu (2019) show that fees become a larger component of miners’ rewards over time, but do not account for technological innovation. Zimmerman (2019) explores the relationship between monetary and speculative usage of a cryptocurrency, settlement capacity, and welfare. Hautsch, Scheuch, and Voigt (2018) study the limits to arbitrage between cryptocurrency exchanges that arise due to blockchain congestion, as described by Makarov and Schoar (2019). Abadi and Brunnermeier (2018) discuss how blockchain technology leads to a trade-off between efficiency and security. None of these papers examine the effect on blockchain space by innovations such as the Lightning Network, SegWit, or Bitcoin Cash.

To our knowledge, ours is the first paper to document the fall in Bitcoin congestion since January 2018, and to find an association with adoption of the Lightning Network. There are only a few papers that study any aspect of the Lightning Network. Bertucci (2020) studies a strategic model of network formation, and shows that competition between nodes prevents the network from becoming highly centralised. Auer (2019) identifies the Lightning Network as an innovation that could ease blockchain congestion, and discusses a tendency toward centralisation. Bartolucci, Caccioli, and Vivo (2019) simulate the Lightning Network using a network percolation model and discuss its feasibility. Ersoy, Roos, and Erkin (2019) and Béres, Seres, and Benczúr (2019) both analyse the profitability of acting as an intermediary in the Lightning Network. Bertucci (2020)

Similarly, there are few papers that look at SegWit and none, to our knowledge, that focus on Bitcoin Cash. Pérez-Solà et al. (2019) discuss the background behind SegWit, and find evidence that is helping to increase blockchain capacity in Bitcoin. Brown, Chiu, and Koepll (2019) model the introduction of SegWit as a positive innovation to the supply of blockchain space, and suggest its abolition would increase miners' revenues. We show that, once adoption of the Lightning Network is accounted for, SegWit adoption actually has no significant impact on blockchain congestion. Lastly, Levine (2019), in an undergraduate thesis, finds empirical evidence that the introduction of SegWit strengthens the relationship between trading volume and returns, but does not examine how it changes the use of blockchain space.

3 Data

We construct measures of mempool congestion using publicly available information from Jochen Hoenicke (<https://jochen-hoenicke.de/queue>). This website provides mempool data at one-minute intervals for several cryptocurrencies, including Bitcoin and Bitcoin Cash. In particular, we collect data on: (i) the number of pending transactions in the mempool (*mempool txn count*); (ii) fees attached to the pending trans-

actions (*mempool txn fees*); and (iii) the proportion of transactions with fees under 10 satoshis per virtual byte (*low fee txns*).³ We use daily data from January 1, 2017 to September 5, 2019. A detailed description of every variable we use, along with sources, is provided in Table 6 in the Appendix.

Hoenicke does not include free transactions — i.e. those with a fee of zero — in his data, even if such transactions are eventually settled on the blockchain. The reason is that it is costless for a vexatious attacker to submit zero-fee transactions to the mempool, so miners often ignore them. Including zero-fee transactions would therefore overstate the actual level of mempool congestion. In any case, Easley, O’Hara, and Basu (2019) show that the number of zero-fee transactions has recently fallen to negligible levels.

The metrics for Bitcoin Cash, SegWit, and the Lightning Network are the independent variables of interest. We call these the ‘innovation indicators’. We investigate whether the drop in mempool congestion and growth in low fee transactions are associated with any of these innovations. As our data begins on January 1, 2017, it includes a period before any of the three innovations were introduced. We obtain data on SegWit from Bitcoin Visuals (<https://bitcoinvisuals.com/chain-tx-block>), who estimate the proportion of Bitcoin transactions that use SegWit in each block. BitMEX Research provide data on the percentage of total daily fees in Bitcoin that use the SegWit protocol (<https://txstats.com/dashboard/db/segwit-usage>). Data on Lightning Networks come from Robtex (<https://hashxp.org/lightning>). This repository contains detailed historical information on all public Lightning nodes (both active and inactive), channels between these nodes (both open and closed), and channel capacity (in bitcoin and USD). In addition, Robtex provides complete details of Bitcoin transactions ex-

³There are 100 million satoshis to a bitcoin. Virtual bytes are a way of accounting for the efficiency of SegWit transaction storage. A virtual byte is equivalent to a physical byte for non-SegWit transactions, and to four physical bytes for SegWit transactions. Hoenicke does not provide fees per physical byte.

ecuted in order to open and close an LN channel. Data on Bitcoin Cash transactions and fees come from Coin Metrics (<https://coinmetrics.io>). Each measure of Bitcoin Cash demand is divided by the equivalent measure of Bitcoin demand, in order to create a relative measure.⁴

We introduce several controls to proxy demand and supply for Bitcoin and blockchain space. Data on all of these controls are obtained from Coin Metrics. *30-day volatility* is the rolling standard deviation of bitcoin returns from each of the past 30 trading days. *1-day price change* is used to control for fluctuations in the price of Bitcoin, and is measured as the rolling difference between days $t - 1$ and $t - 2$ in log bitcoin price at 0000 hours UTC.

We include a supply measure, called *mining intensity*. It estimates the average rate of block creation per unit time. Miners create new blocks by picking up transactions from the mempool, and attempting to solve a complex mathematical puzzle. When a miner solves the puzzle, a new block is created, to which the selected mempool transactions are added. The successful miner is rewarded with the fees attached to these transactions, plus some newly minted bitcoins.⁵ The new block is then added to the blockchain. We define mining intensity as the total computational power used by miners to solve the puzzle (called ‘hash rate’), divided by its difficulty.⁶

Table 1 shows summary statistics for the variables. In the following section, we briefly describe each of the three innovations and discuss the extent of their adoption over the time period we study.

⁴We carry out this normalisation because we are interested in the extent to which Bitcoin Cash substitutes for Bitcoin. It is intended to avoid spurious results that may arise if there are factors that drive demand for both Bitcoin and Bitcoin Cash in the same direction. For example, if market sentiment becomes improves toward all cryptocurrencies, then demand for both will go up. A regression may then erroneously suggest that increased demand for Bitcoin Cash leads to worse congestion for Bitcoin. The normalised measure is a better indicator of relative demand.

⁵At the time of writing, mining a block earns new 12.5 bitcoins, worth approximately \$91,000.

⁶Bitcoin difficulty is adjusted every 2,016 blocks to target an average block creation rate of roughly one every ten minutes. For example, if miners increase their hash rate, then at the next adjustment, difficulty will increase. As hash rate tends to increase over time, the actual average block creation rate is usually faster than one every ten minutes.

Table 1: Summary statistics.

Daily data from January 1, 2017 to September 5, 2019. See Table 6 for variable definitions and sources of data.

	count	mean	std dev	min	median	max
Mempool txn count	972	23,042	40,619	92	5,731	252,750
Mempool txn fees (USD)	965	106,180	440,206	39	3,008	4,750,619
Low fee txns (%)	965	53.45	28.30	0	52.04	95.99
BCH/BTC txns	972	0.14	0.58	0	0.07	10.26
BCH/BTC fees	972	0.21	0.39	0	0.07	8.46
SegWit txns (%)	972	20.72	15.48	0	27.61	46.80
SegWit txns by fee (%)	965	19.75	15.98	0	18.31	51.57
Lightning Network channels	972	12,671	15,374	0	7,575	44,087
Lightning Network capacity (USD)	972	2,766,535	4,080,066	0	205,388	11,794,337
30-day volatility	972	4.16	1.54	1.10	4.03	8.07
1-day price change	965	0.002	0.044	-0.207	0.003	0.225
Mining intensity	972	7.49	0.82	3.98	7.51	9.79

4 Adoption of the technological innovations

Bitcoin Cash (BCH) launched on August 1, 2017 as a *hard fork* of Bitcoin. A hard fork occurs when a new version of the core code is released and then adopted by a subset of miners of the original blockchain. Blocks created under the old code continue to be added to the Bitcoin blockchain, while blocks created under the new code are part of a new blockchain. In this way, a new currency is ‘forked’ from the existing one without replacing it. The developers programmed Bitcoin Cash to have greater settlement capacity than Bitcoin, in order to function better as a means of payment. BCH blocks can hold up to 8MB of data, compared to 1MB for Bitcoin.⁷

⁷For more on Bitcoin Cash, see <https://www.bitcoincash.org/index.html>. On November 15, 2018, Bitcoin Cash was itself hard-forked to create yet another cryptocurrency, *Bitcoin Satoshi’s Vision* (SV), with capacity of up to 128MB per block. See <https://bitcoinsv.io>. We do not analyse Bitcoin SV or other forked cryptocurrencies, because they tend to have negligible market capitalisation relative to Bitcoin and Bitcoin Cash.

SegWit (an acronym for ‘Segregated Witness’) is an on-chain scaling solution that changes the way transactions are stored.⁸ It was activated on August 23, 2017 via a *soft fork*, which is a consensual update of certain attributes of the blockchain. The update did not change the block size, but improved the efficiency of transaction storage, so that a block can potentially store up to four times as many transactions as before. It preserved other characteristics and did not split the blockchain in two.

The Lightning Network (LN) went live at the beginning of 2018. Rather than aiming to increase blockchain capacity, LN is an *off-chain* solution. It is a secondary transaction layer that operates outside of the Bitcoin blockchain (Poon and Dryja 2016). Two Bitcoin users can open an LN channel through which they can make payments to one another. Once the channel is closed, only the net amount needs to be settled on-chain, as a single payment. This compresses the number of on-chain transactions required to just two (one to open an LN channel and another to close it), allowing the system to handle a much larger number of payments. LN users have to lock bitcoin into the channels in order to collateralise their positions.

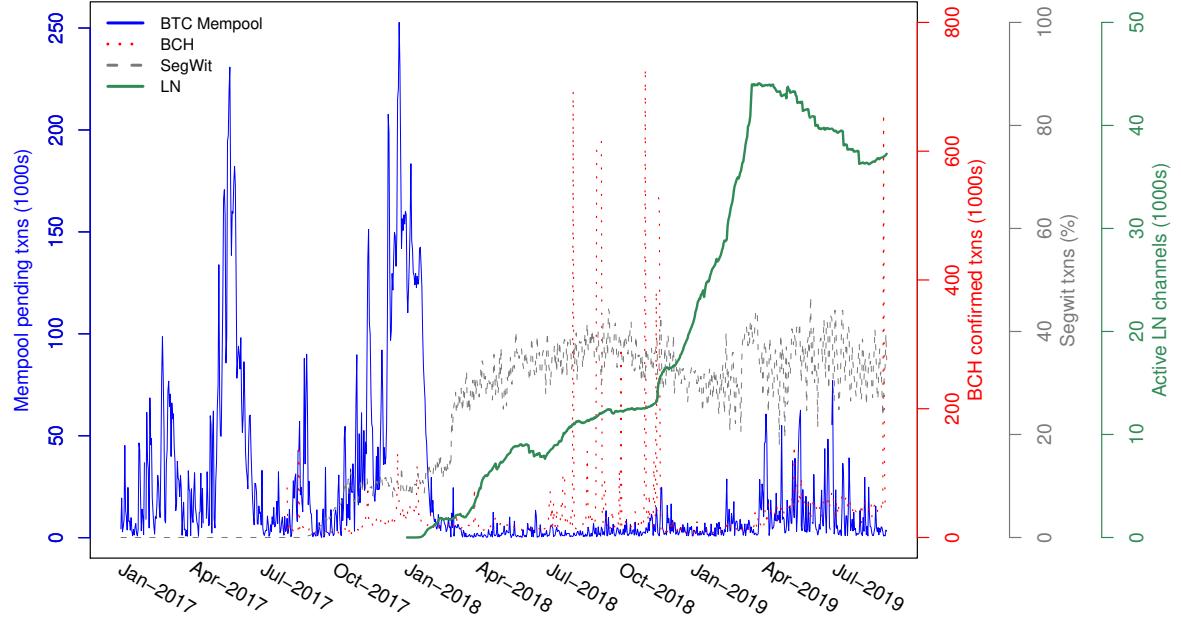
Each of these three innovations has the potential to reduce congestion in the Bitcoin mempool. Figure 3 plots the number of unconfirmed mempool transactions (blue line) against usage of the three technologies. The red line shows the number of Bitcoin Cash transactions, the grey line the proportion of Bitcoin transactions that are made via SegWit, and the green line the number of active Lightning Network channels. There was a dramatic reduction in Bitcoin mempool congestion at the beginning of 2018 when the speculative bubble burst. Congestion has remained relatively low since then, but picked up slightly in mid-2019.

Similarly, Figure 4 plots transactions weighted by fees. The blue line shows that the total fees attached to payments waiting in the Bitcoin mempool has fallen since 2017, suggesting either lower demand or greater supply of settlement capacity. Over this period, the fee-weighted proportion of SegWit transactions has risen (grey line), as has the total value of bitcoin used to collateralise Lightning channels (green line), suggesting that users are paying more to use these services. The red line shows the total fees paid for settled Bitcoin Cash payments. This peaks in 2017 at the same time as Bitcoin mempool fees, which is perhaps indicative of the cryptoasset boom that year.

⁸See <https://bitcoincore.org/en/2016/01/26/segwit-benefits>.

Figure 3: Bitcoin mempool size and adoption of innovations.

Daily data from January 1, 2017 to September 5, 2019. *BTC* stands for Bitcoin, *BCH* for Bitcoin Cash, and *LN* for Lightning Network.

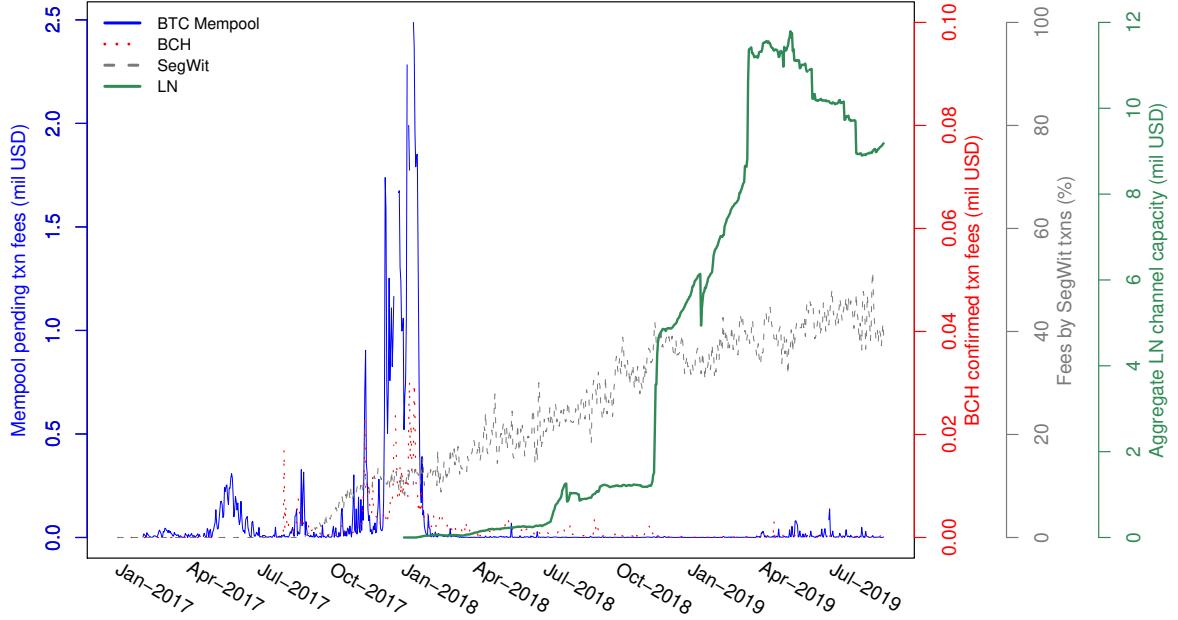


Figures 3 and 4 suggest the potential presence of time trends, since mempool congestion has fallen since the beginning of 2018, around the same time as the introduction of the technological innovations. Augmented Dickey-Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests confirm that the three measures of mempool characteristics — *mempool txn count*, *mempool txn fees (USD)*, and *low fee txns (%)* — are non-stationary, as are several of our independent variables.⁹ We therefore follow the Box–Jenkins method and take first differences. ADF and KPSS tests suggest these first-differenced variables are stationary.

⁹The indicators for SegWit and Lightning Network adoption, along with 30-day volatility and mining intensity, are all non-stationary.

Figure 4: Bitcoin mempool fees USD versus innovations.

Daily data from January 1, 2017 to September 5, 2019. *BTC* stands for Bitcoin, *BCH* for Bitcoin Cash, *SegWit* for Segregated Witness, and *LN* for Lightning Network.



5 Empirical results

We test for an association between adoption of the technological innovations and Bitcoin mempool congestion. We use autoregressive integrated moving average (ARIMA) models. These postulate that the variable of interest can be written as a function of its past values of the parameter of interest, past forecast errors, and other predictors. An ARIMA model with parameters (p, d, q) tests the following specification:

$$y_t^d = c + \sum_{i=1}^p \phi_i y_{t-i}^d + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + X_t^d \beta + \varepsilon_t, \quad (1)$$

where c is a constant term, y^d is the variable of interest expressed after taking d differences, X_t^d is a vector of the d -differenced independent variables, and ε_t is a residual term. The parameter p is the number of lags of the variable of interest included on the right-hand side, d is the number of differences taken, and q is the length of window for the moving average of historic residual terms. For each specification, we estimate the parameters (p, d, q) .¹⁰

Table 2 reports the effect of each innovation on mempool txn count. This is the number of transactions in the mempool waiting to be confirmed and added to the blockchain. We run regressions on each of the three innovations, with and without the control variables. We also run regressions with all three innovations included. Greater usage of the Lightning Network is associated with a lower mempool count, while Bitcoin Cash and SegWit have no significant impact. None of the supply and demand controls have a significant impact on mempool size. The last three rows in Table 2 report results of the Portmanteau Q-test and Durbin-Watson test, which are consistent with no autocorrelation in the residuals. We estimate, using model (8) in Table 2, that an increase of one standard deviation in the number of Lightning Network channels is associated with a decrease in the mempool count of approximately 0.31 standard deviations.

Table 3 reports results for ARIMA regressions of the three innovations on aggregate transaction fees in the Bitcoin mempool. For these models, we have used fee-based innovation indicators. For example, we use the proportion of SegWit transactions by fee, rather than by number. As before, Lightning Network capacity is associated with a reduction of transaction fees in the mempool, although the significance of these effects is lower. We estimate from model (8) in Table 3 that an increase of one standard deviation in the dollar-denominated capacity of Lightning Networks is associated with a drop of 0.11 standard deviations in aggregate dollar-denominated transactions fees in the Bitcoin mempool.

¹⁰We use the `auto.arima` function in R. This employs the Hyndman-Khandakar algorithm (Hyndman and Khandakar 2008), which uses stepwise search to identify the parameters (p, d, q) with the lowest Akaike information criterion.

Table 2: Impact of technological innovations on Bitcoin mempool count.

ARIMA regression results of the three innovation measures on mempool transaction count (log). In all eight models, the parameters selected by the Hyndman-Khandakar algorithm are the same: $p = 6$ lagged terms included for the dependent variable, $d = 1$ difference taken, and $q = 2$ length of window for the moving average of historical residual terms. The data is from January 1, 2017 to September 5, 2019. See Table 6 for variable definitions and data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ BCH/BTC	-0.001	-0.000					-0.001	-0.001
txns	(0.000)	(0.001)					(0.000)	(0.001)
Δ SegWit txns (%)			0.016 (0.012)	0.017 (0.012)			0.016 (0.012)	0.017 (0.012)
Δ LN channels (log)					-0.257*** (0.074)	-0.273*** (0.075)	-0.262*** (0.076)	-0.260*** (0.078)
Δ 30-day volatility		-0.020 (0.081)		-0.014 (0.082)		-0.020 (0.081)		-0.023 (0.082)
Δ 1-day price change		-0.672 (0.630)		-0.681 (0.629)		-0.664 (0.630)		-0.770 (0.620)
Δ Mining intensity		0.037 (0.049)		0.047 (0.049)		0.039 (0.049)		0.037 (0.049)
Constant	-0.004 (0.009)	-0.004 (0.008)	-0.004 (0.009)	-0.004 (0.008)	-0.001 (0.009)	-0.001 (0.009)	-0.001 (0.009)	-0.002 (0.009)
Observations	965	965	965	965	965	965	965	965
AIC	2589	2598	2588	2597	2588	2597	2589	2590
BIC	2643	2667	2642	2665	2642	2665	2652	2667
Q	3.237	10.462	3.082	10.123	4.000	10.099	3.821	3.421
$p(Q)$	0.975	0.401	0.979	0.430	0.947	0.432	0.955	0.970
Durbin-Watson	1.977	1.989	1.977	1.989	1.976	1.988	1.977	1.974

As in Table 2, demand for Bitcoin Cash has no significant impact on Bitcoin mempool transaction fees. SegWit adoption has a slightly positive impact, which is the opposite of what might be expected.¹¹

¹¹ One explanation is that the incentive to use SegWit may be strong only for transactions that take up a large amount of space. During periods when there is more demand for bulkier transactions, both the proportion of SegWit transactions and mempool congestion may go up. Moreover, many cryptocurrency exchanges have been reluctant to make their infrastructures SegWit-compliant due to the high costs involved. See the discussion at <https://tinyurl.com/sjfs7d8>.

Table 3: Impact of technological innovations on Bitcoin mempool fees.

ARIMA regression results of the three innovation measures on mempool fees (USD log). In all eight models, the parameters selected by the Hyndman-Khandakar algorithm are the same: $p = 6$ lagged terms included for the dependent variable, $d = 1$ difference taken, and $q = 1$ length of window for the moving average of historical residual terms. The data is from January 1, 2017 to September 5, 2019. See Table 6 for variable definitions and data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ BCH/BTC fees	-0.031 (0.112)	-0.025 (0.109)					-0.038 (0.116)	-0.032 (0.113)
Δ SegWit txns by fee (%)			0.039* (0.023)	0.039* (0.023)			0.040* (0.023)	0.041* (0.023)
Δ LN capacity (USD log)					-0.189** (0.095)	-0.195** (0.095)	-0.199** (0.094)	-0.205** (0.093)
Δ 30-day volatility	0.088 (0.108)		0.077 (0.107)			0.090 (0.107)		0.077 (0.106)
Δ 1-day price change	0.651 (0.445)		0.669 (0.443)			0.657 (0.445)		0.679 (0.442)
Δ Mining intensity	0.034 (0.059)		0.036 (0.059)			0.034 (0.059)		0.035 (0.059)
Constant	-0.003 (0.010)	-0.003 (0.010)	-0.005 (0.010)	-0.005 (0.010)	0.000 (0.010)	0.000 (0.010)	-0.001 (0.010)	-0.001 (0.010)
Observations	965	965	965	965	965	965	965	965
AIC	2907	2909	2903	2905	2905	2907	2905	2907
BIC	2950	2967	2947	2963	2949	2965	2958	2975
Q	15.957	15.375	15.658	15.077	14.678	14.107	14.575	13.989
$p(Q)$	0.101	0.119	0.110	0.129	0.144	0.168	0.148	0.173
Durbin-Watson	2.025	2.022	2.022	2.020	2.021	2.018	2.019	2.016

Finally, we investigate how the three innovations affect the proportion of mempool transactions that have low fees attached (defined as less than 10 satoshis per virtual byte). Table 4 shows the results. Increased usage of the Lightning Network is significantly associated with an increase in low fee transactions, in line with predictions. Settlement via the Lightning Network reduces the number of transactions that need to be settled on-chain, leading to a drop in the fees that users need to offer. Estimates in

model (8) suggest that a one standard deviation rise in the number of Lightning Network channels is associated with a growth in the percentage of low fee transactions of 0.37 standard deviations. Bitcoin Cash usage has no significant impact on low fee transactions in the mempool.

SegWit usage appears to be negatively related to the proportion of low fee transactions, which is again surprising. It is in line with the findings of Table 3, although the significance is higher. One additional explanation (aside from those in footnote 11) is that our definition of low fee transactions relates to virtual bytes, rather than physical bytes. As SegWit transactions require fewer virtual bytes than non-SegWit transactions, the fee per virtual byte may not actually be lower. The fee per physical byte should be lower, but unfortunately we do not have access to such data.

Overall, these results suggest that increased use of the Lightning Network is associated with a significant reduction in mempool congestion, but the other innovation indicators and control variables do not. As there is no theoretical upper limit on Lightning Network usage, there is potential for still further reductions in congestion in the future.

6 Evolution of the Lightning Network

6.1 Centralisation

We examine how the shape of the Lightning Network has changed over time. Figures 5 and 6 provide snapshots of the structure of the Lightning Network on April 1, 2018 and August 31, 2019 respectively. Each black node depicts a LN user, and each red line represents a LN channel between those nodes. The thickness of a line represents the channel capacity; that is, the number of bitcoin pledged to that channel. The size of a node represents its relative importance in the network, measured by eigenvector centrality.¹²

¹²Eigenvector centrality is a relative measure of the importance of a node within a network. A node has a high score if it links to other nodes with high scores. Eigenvector centrality is normalised to lie between 0 and 1, with the most important node assigned a score of 1. The layout of each figure is drawn using the force-directed algorithm of Fruchterman and Reingold (1991).

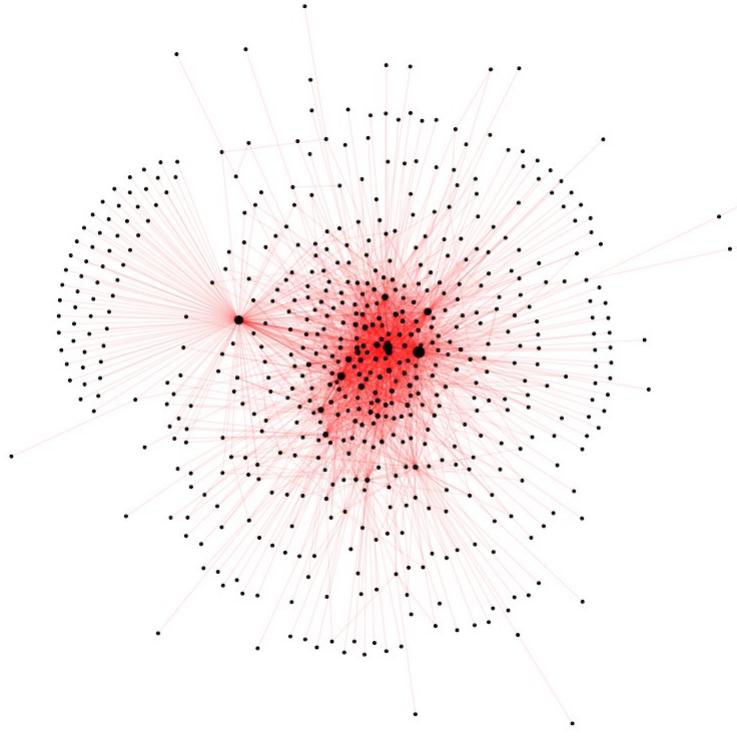
Table 4: Impact of technological innovations on Bitcoin mempool fee distribution.

ARIMA regression results of the three innovation measures on low fee transactions (%). In all eight models, the parameters selected by the Hyndman-Khandakar algorithm are the same: $p = 6$ lagged terms included for the dependent variable, $d = 1$ difference taken, and $q = 2$ length of window for the moving average of historical residual terms. The data is from January 1, 2017 to September 5, 2019. See Table 6 for variable definitions and data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ BCH/BTC	-0.000	-0.000					-0.000	-0.000
txns	(0.000)	(0.000)					(0.000)	(0.000)
Δ Segwit txns (%)			-0.017*** (0.004)	-0.015*** (0.004)			-0.017*** (0.004)	-0.016*** (0.004)
Δ LN channels (log)					0.186*** (0.035)	0.178*** (0.036)	0.189*** (0.035)	0.180*** (0.036)
Δ 30-day volatility	-0.031 (0.027)		-0.037 (0.027)		-0.024 (0.025)		-0.030 (0.024)	
Δ 1-day price change	0.046 (0.224)		0.051 (0.223)		0.047 (0.224)		0.058 (0.222)	
Δ Mining intensity	0.041*** (0.015)		0.036** (0.015)		0.041*** (0.015)		0.034** (0.015)	
Constant	0.000 (0.002)	0.000 (0.002)	0.001 (0.002)	0.001 (0.002)	-0.002 (0.002)	-0.002 (0.002)	-0.002 (0.002)	-0.001 (0.002)
Observations	965	965	965	965	965	965	965	965
AIC	710	690	691	675	694	676	678	661
BIC	764	759	745	743	748	744	741	739
Q	4.359	2.868	2.816	1.886	2.881	2.284	2.081	1.389
$p(Q)$	0.930	0.984	0.985	0.997	0.984	0.994	0.996	0.999
Durbin-Watson	1.955	1.982	1.955	1.987	1.946	1.981	1.950	1.986

The figures suggest that the Lightning Network has become more centralised between April 1, 2018 and August 31, 2019, characterised by a small number of highly connected users. Table 5 compares network statistics between these two dates. The number of participants in the LN has risen five-fold from 1,145 to 5,699. However, the average degree — the mean number of nodes each node connects to — has increased only from 7.7 to 11.4. Therefore, the connectivity — the probability that any two randomly selected nodes are connected — has fallen from 0.7% to 0.2%. The mean eigenvector

Figure 5: Lightning Network structure on April 1, 2018. Network diagram constructed on Gephi using the Fruchterman-Reingold layout algorithm. Scale of nodes and links is relative and not drawn to same scale as Figure 6.



centrality of the five most-connected nodes has risen, suggesting that the network is increasingly reliant on a few core nodes. These statistics are consistent with a network in which a few core nodes become more important as intermediaries, while the rest remain peripheral.

Table 5: Network statistics of the Lightning Network.

Network Statistic	Apr 1, 2018	Aug 31, 2019
Nodes (LN users)	1145	5699
Edges (LN channels)	4401	32594
Average degree	7.687	11.438
Average connectivity	0.007	0.002
Average path length	3.107	3.297
Network diameter	8	11
Top 5 mean eigenvector centrality	0.690	0.876

Figure 6: Lightning Network structure on August 31, 2019. Network diagram constructed on Gephi using the Fruchterman-Reingold layout algorithm. Scale of nodes and links is relative and not drawn to same scale as Figure 5.

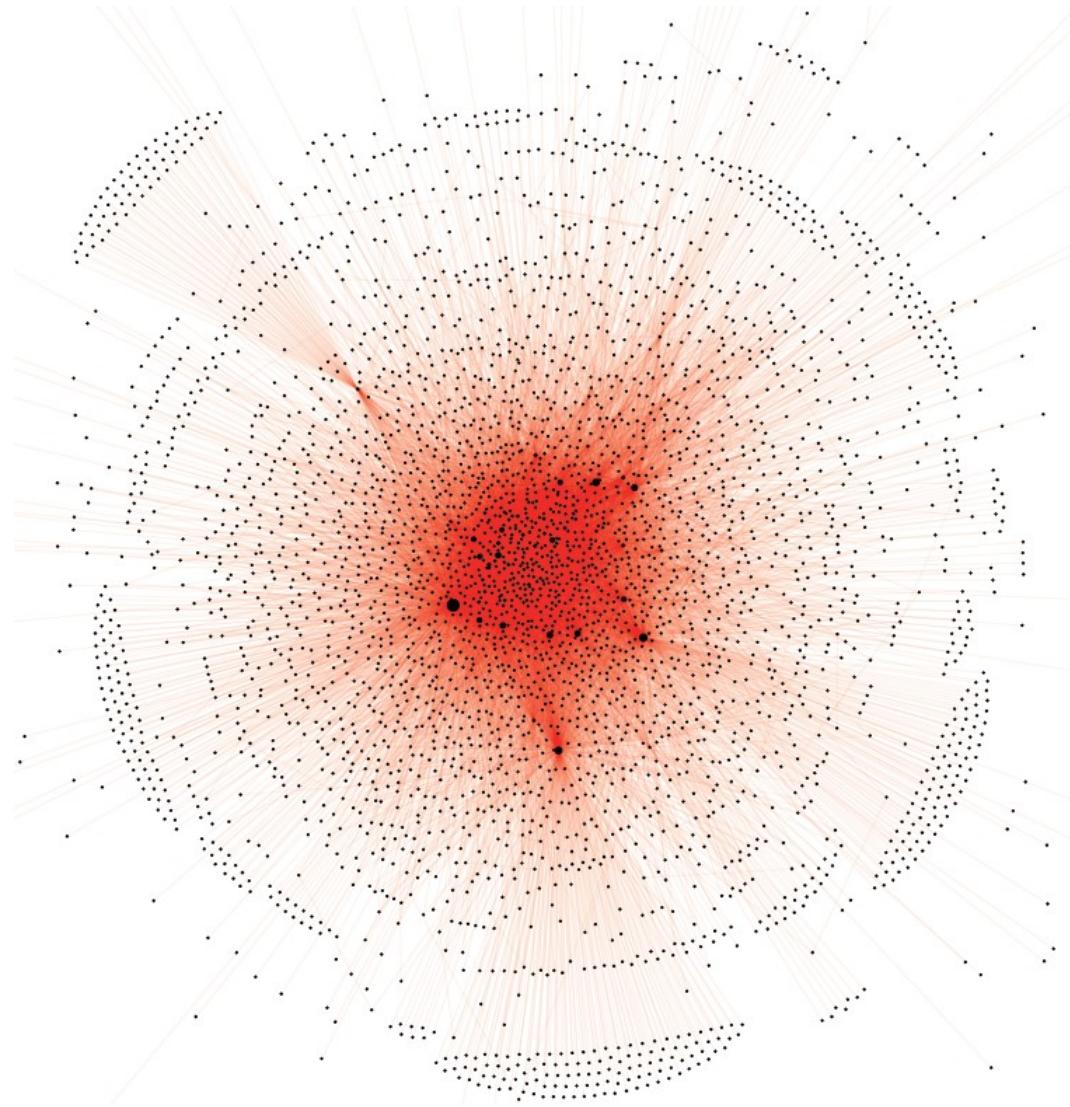
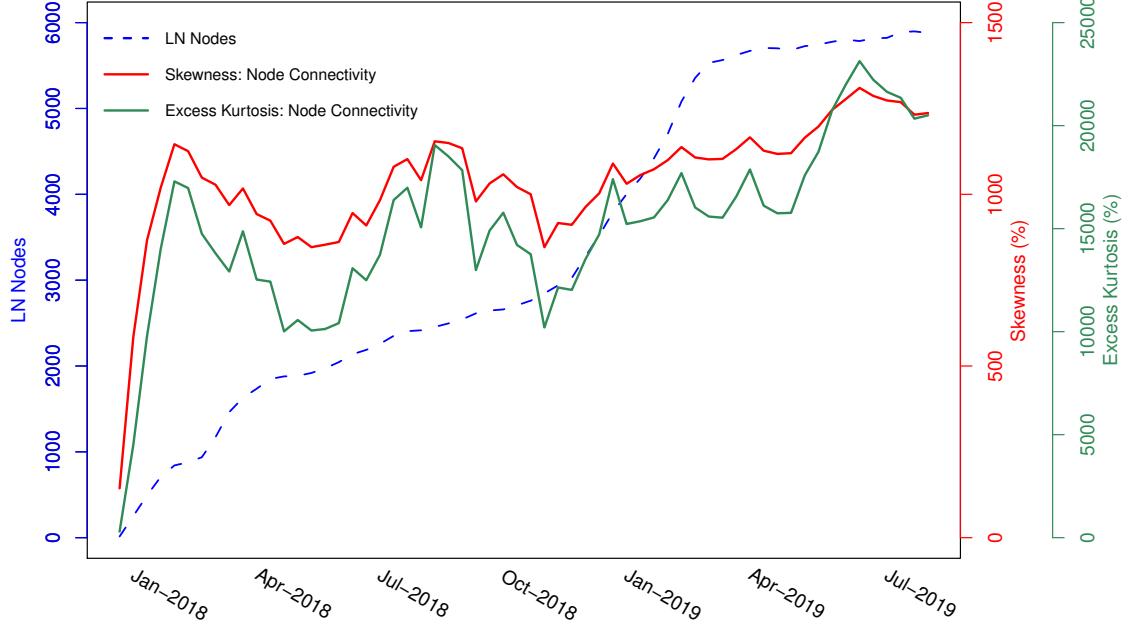


Figure 7 plots the skewness and excess kurtosis of the distribution of connectivity over time, along with the number of nodes in the network. The connectivity distribution has positive skewness and excess kurtosis. This means there is a fat tail of Lightning nodes that are highly connected, and a large number of sparsely connected nodes. Both skewness and excess kurtosis tend to rise over time, consistent with increased centralisation.

Figure 7: Connectivity among Lightning nodes: distributional properties.

Values are computed once every 10 days over a rolling 30-day window.



6.2 Economic factors driving centralisation

In this section, we argue that the increased centralisation of the Lightning Network permits it to be used at lower cost. Usually, net settlement between two counterparties creates default risk. The LN avoids this by requiring both counterparties to collateralise their channels. This is done by locking bitcoin into a smart contract when a new channel is opened. An LN counterparty is not permitted to take a net debit position in excess of this collateral. There is an opportunity cost to opening a Lightning channel, since bitcoin have to be locked up while it is open. The cost may be the missed opportunity to use Bitcoin as money, or to make short-term trades in the cryptocurrency market.

Centralisation of the Lightning Network reflects increased use of *multi-hop channels*. Suppose Alice wishes to make a bitcoin payment to Bob. Rather than open a direct Lightning channel with Bob, Alice can instead send the coins via an intermediary, Xavier. This can be done so long as Alice and Bob both have open channels with Xavier, and

there is sufficient capacity in the channels along this path to support the payment. In this example, there is a channel from Alice to Bob with two hops. There is no theoretical limit to the number of hops that a payment can take to reach its destination. For example, Alice’s payment might pass through Xavier, Yvette and Zebedee, before reaching Bob.

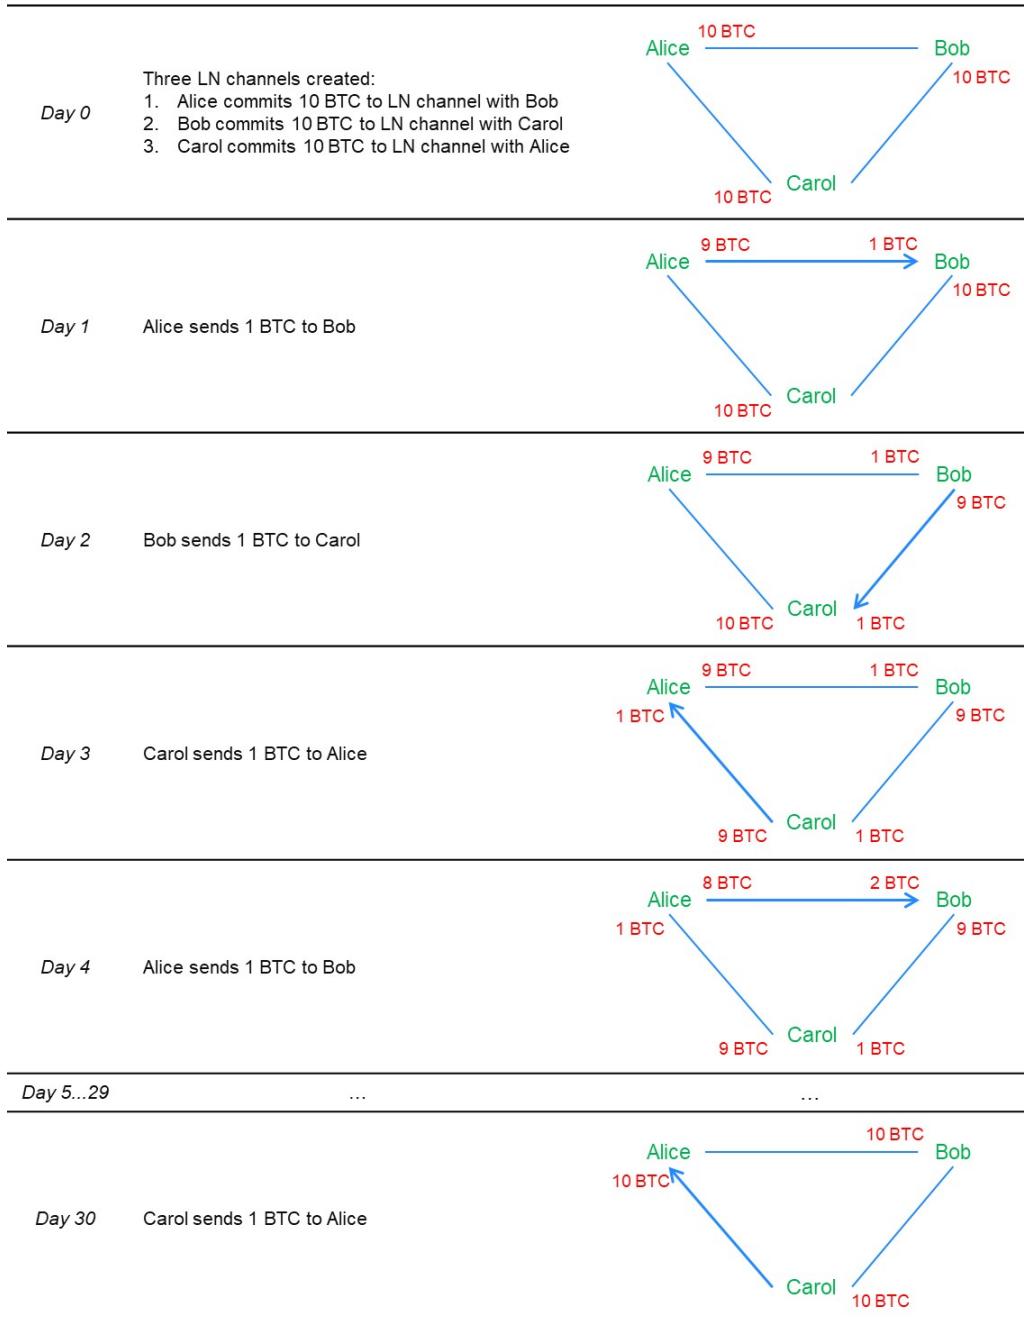
Multi-hop channels can reduce the aggregate amount of bitcoin that is needed to be posted as collateral to the network. This can be illustrated with a simple example. Suppose Alice knows she has to pay Bob 1 bitcoin every third day for the next 30 days; i.e. she pays him on days 1, 4, 7, . . . , 28. Similarly, Bob pays Carol 1 bitcoin on days 2, 5, 8, . . . , 29, and Carol pays Alice 1 bitcoin on days 3, 6, 9, . . . , 30. Suppose further that all LN channels have to be opened on day zero and cannot be revised or closed before day 30. Let r be the opportunity cost of locking up a bitcoin for 30 days.

Let us first consider the cost of the Lightning Network with direct channels. Three channels are opened, one for each pair of counterparties, and each must have 10 bitcoins of collateral. The total system-wide cost is $30r$. Figure 8 shows how settlement occurs with direct channels.¹³

Now suppose that, instead of opening direct channels with one another, Alice, Bob, and Carol all open channels with Xavier and direct their payments via him. On days 1, 4, 7, etc, Alice has a debit position with Xavier of 1, so must pledge 1 bitcoin. On these days, Xavier has a debit position with Bob of 1 and, on days 2, 5, 8, etc, Xavier has a

¹³Clearly, we can do better if Alice can close the channel partway through and reopen it using the money paid to her by Carol, and so forth. We rule that out in this example. We assume a high cost of closing and reopening a Lightning channel. Two on-chain transaction need to be made, and so fees must be paid to miners.

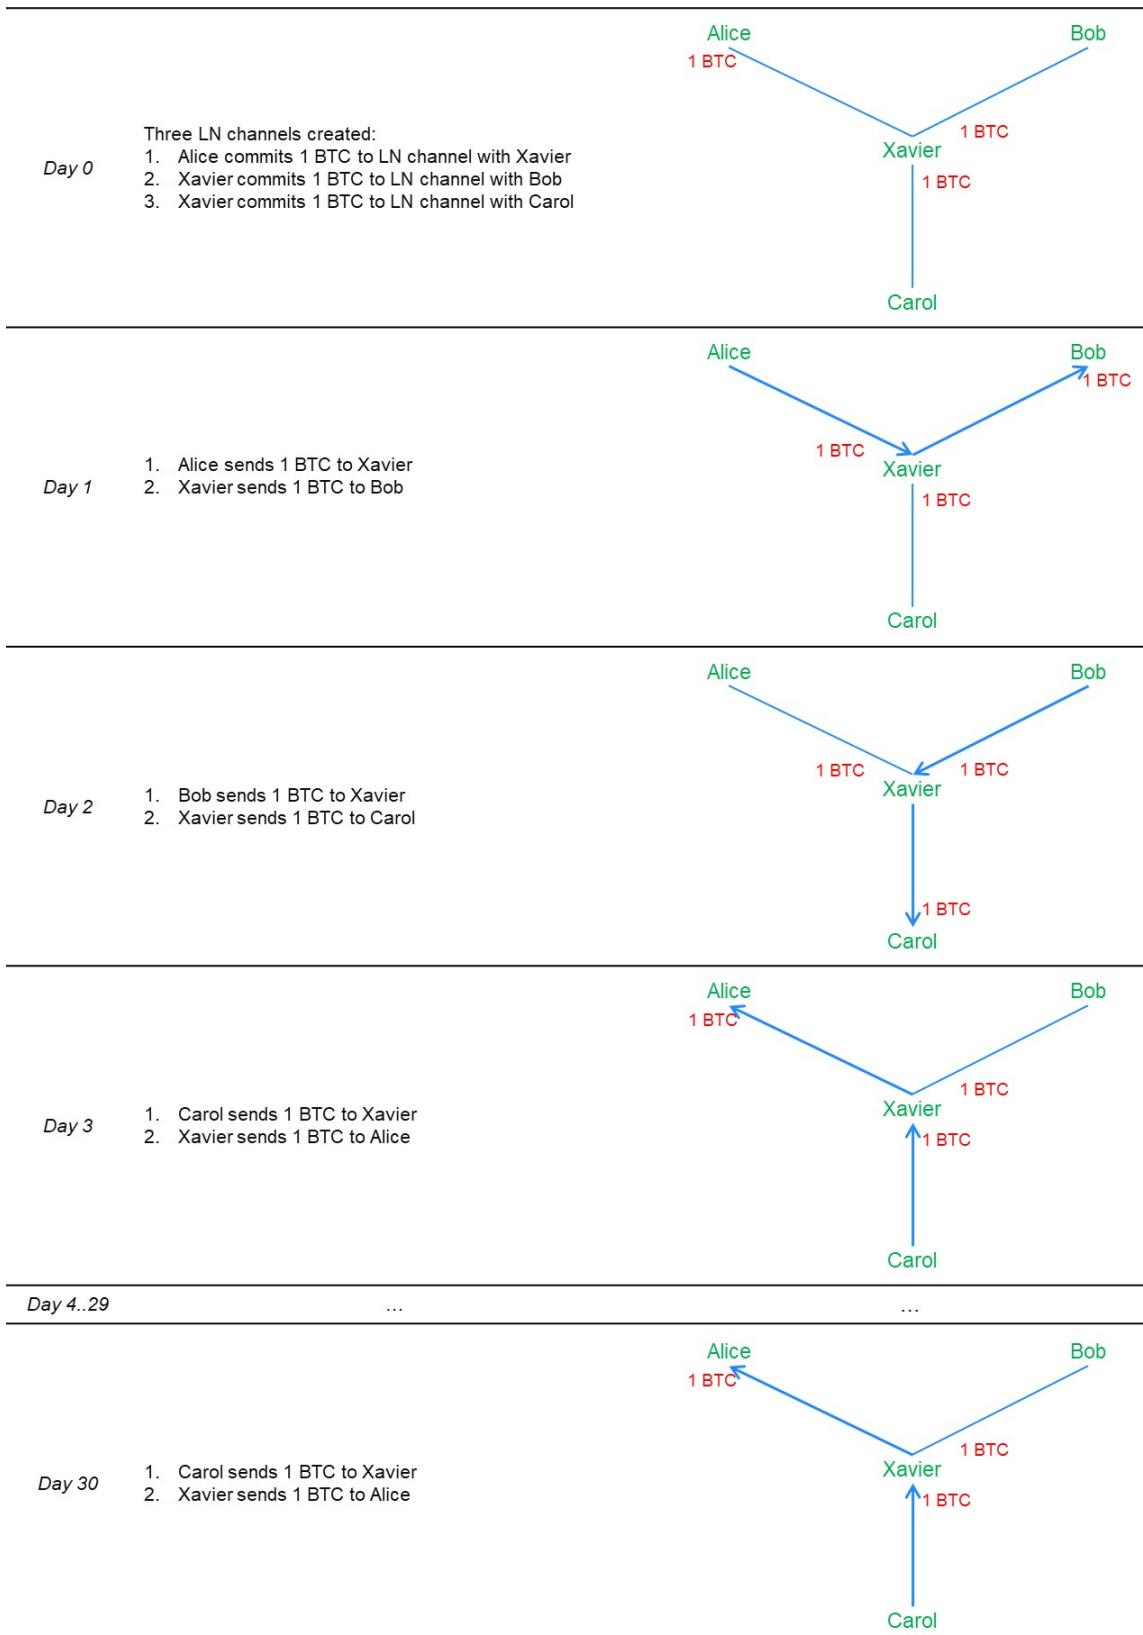
Figure 8: Payments through direct LN channels. Net debit positions in red.



debit position of 1 with Carol. Thus Xavier must pledge 1 bitcoin to each of his channels with Bob and Carol. For their part, Bob and Carol always receive before they pay, so they do not need to pledge anything. The total amount to be pledged is 3 bitcoins. Figure 9 shows how this works.¹⁴

¹⁴Note that the Lightning Network always requires more bitcoin to be locked up than if the same payments were made on-chain. If Alice, Bob and Carol were to make payments directly on-chain, then Alice would just need to contribute a single bitcoin, which would facilitate all subsequent payments. This implies a system-wide cost of r . But there would be additional costs: in particular, the participants would have to pay fees to miners, and may face delayed settlement if the mempool is congested. Users have to trade off these costs when decided whether to use the Lightning Network or not.

Figure 9: Payments through LN intermediaries. Net debit positions in red.



In this example, the system-wide cost of operating multi-hop channels is $3r$, a small fraction of the cost of using direct channels. The total value to society of Xavier’s intermediary service is $27r$. These substantial economies of scale demonstrate the incentive to use multi-hop channels, and help explain the increased centralisation described in Table 5 and Figure 7.

Competitive forces may prevent the Lightning Network from becoming completely centralised, with one node intermediating all payments. This is because intermediaries can elect to charge fees for providing channels. Partly these fees serve to compensate an intermediary for the bitcoin they have to lock up, but they can also provide a profit motive.¹⁵ If Xavier has monopoly power, he could potentially make a profit of up to $27r$, the total value of his service to society. He can thus extract significant rents from Alice, Bob, and Carol. However, other agents, such as Yvette and Zebedee, can enter the network and try to compete with Xavier, reducing the profit. But there are barriers to entry: for example, if Xavier has already established channels with Alice’s counterparties but Yvette has not, he may be able to provide Alice with intermediary services at a lower cost.

This discussion suggests that centralisation may arise endogenously in the Lightning Network, because agents prefer to connect to counterparties that can route payments at lowest cost. But competitive forces may prevent the network from becoming too centralised, with only a few core nodes routing all the payments. The actual extent of centralisation is likely to depend on the trade-off between the benefit of economies of scale (increasing in the cost of locking up bitcoin) and the cost of centralisation (increasing in the cost of entry). We will model this trade-off in follow-up work.

¹⁵The fees charged by a Lightning intermediary are not to be confused with those charged by miners for confirming transactions on the blockchain.

7 Welfare

The Lightning Network has positive consequences for welfare. First, as the Lightning Network makes Bitcoin a better payments system, users are better off. Their transactions settle more quickly and more cheaply (Zimmerman 2019). Second, as fewer transactions need to be recorded on the blockchain, less memory and energy is needed to run a Bitcoin node. This reduces the cost of maintaining the blockchain, allowing more nodes to participate and making the system more secure against a double-spending attack (Budish 2018).

Finally, the Lightning Network may reduce the energy consumed by Bitcoin miners, generating positive externalities for the broader society. Greater use of the Lightning Network reduces the fees paid to miners, which in turn lowers their incentive to devote computing power (hash rate) to the network. In the medium term, this has no negative effect on the supply of new blocks, because the difficulty adjusts accordingly. But it does mean lower energy usage and thus less generation of pollutants. The total energy consumption of Bitcoin miners is substantive, so the benefits could potentially be large.¹⁶ However, this benefit is likely to be realised only in the long term because, currently, fees comprise a very small part of miners' revenue. Fees are expected to grow in importance as block rewards continue to decline over time (Easley, O'Hara, and Basu 2019).

There is a potential negative effect on welfare too. Faced with reduced revenue, the miners who face highest costs may exit the market. This could allow relatively few miners to control the blockchain and manipulate it. However, this risk is offset somewhat by the reduction in memory space needed to record new blocks, reducing the cost of running a node.

¹⁶At the time of writing, Bitcoin is estimated to generate as much carbon dioxide as the entire nation of Denmark. See <https://digiconomist.net/bitcoin-energy-consumption>.

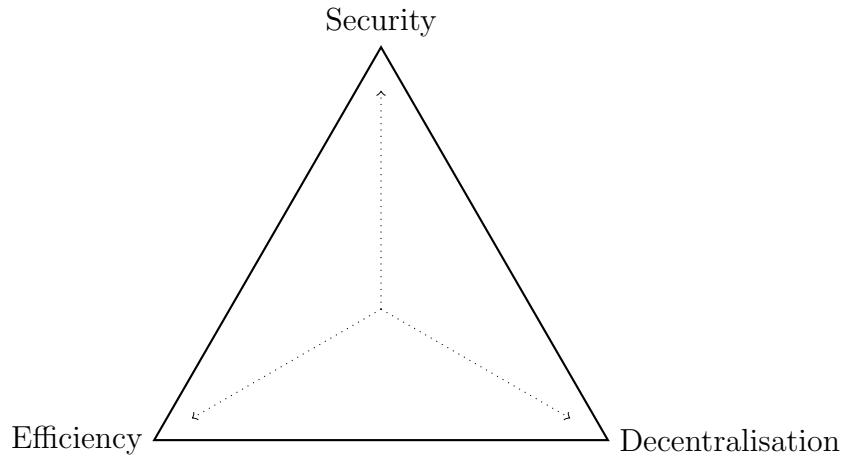
8 Conclusions

We show that usage of the Lightning Network is associated with reduced mempool congestion in Bitcoin, and lower fees. This suggests that the netting benefits of the Lightning Network can help Bitcoin function better as a means of payment. Data on actual Bitcoin usage are not available, so we cannot say for sure whether Bitcoin is being increasingly used as money. However, we can say that the Lightning Network loosens a key technological constraint by allowing payments to be settled more quickly. It may also reduce barriers to arbitrage across cryptocurrency exchanges, as identified by Makarov and Schoar (2019) and Hautsch, Scheuch, and Voigt (2018), thereby improving market liquidity.

Economies of scale in the Lightning Network arise when there are multi-hop channels, with users sending each other payments via intermediary nodes. Because of this, the network has become increasingly centralised. Centralisation allows intermediaries to extract rent, which reduces the economic benefit to users. Competition will prevent the network from becoming completely centralised, and thus there is a limit to the extent to which the Lightning Network can reduce mempool congestion.

The blockchain-based structure of Bitcoin gives rise to a trilemma (Figure 10). A currency can achieve two of decentralisation, security, and payments efficiency, but not all three. For example, fiat currencies are secure and efficient, but are dependent on a central bank. Bitcoin is currently a decentralised and secure currency, but is not an efficient means of payment. The solid black line in Figure 10 represents a technological frontier. Adoption of the Lightning Network pushes this frontier out: it can allow Bitcoin to achieve greater efficiency, while remaining decentralised and secure. There is, however, a trade-off: while the blockchain remains decentralised, economies of scale cause the Lightning Network itself to become increasingly centralised and reliant on a few large nodes.

Figure 10: Trilemma of decentralised money. A currency can achieve two of decentralisation, security, and payments efficiency, but not all three. The solid line represents a technological frontier, which can be expanded by innovations such as the Lightning Network.



The Lightning Network provides a technological solution to allow a blockchain-based decentralised cryptocurrency to achieve scalability. While this paper has focused on Bitcoin, the same technology could allow other currencies to be widely used, secure, and decentralised. For example, the Libra Association has proposed a currency that would be globally used and, ultimately, operate on a blockchain (Libra Association 2019). Lightning technology provides an effective way to achieve such ambitious objectives.

9 References

Abadi, J. and M. Brunnermeier (2018). “Blockchain economics”. *NBER working paper* 25407.

Athey, S., I. Parashkevov, V. Sarukkai, and J. Xia (2016). “Bitcoin pricing, adoption and usage: theory and evidence”. *SSRN* 2826674.

Auer, R. (2019). “Beyond the doomsday economics of proof-of-work in cryptocurrencies”. *BIS working paper* 765. URL: <https://www.bis.org/publ/work765.htm>.

Bartolucci, S., F. Caccioli, and P. Vivo (2019). “A percolation model for the emergence of the Bitcoin Lightning Network”. URL: <https://arxiv.org/pdf/1912.03556.pdf>.

Béres, F., I. A. Seres, and A. A. Benczúr (2019). “A cryptoeconomic traffic analysis of Bitcoin’s Lightning Network”. URL: <https://arxiv.org/pdf/1911.09432.pdf>.

Bertucci, L. (2020). “Incentives on the Lightning Network : a blockchain-based payment network”. *SSRN* 3540581.

Biais, B., C. Bisière, M. Bouvard, C. Casamatta, and A. Menkveld (2018). “Equilibrium bitcoin pricing”. *SSRN* 3261063.

Bolt, W. and M. R. C. van Oordt (2019). “On the value of virtual currencies”. *Journal of Money, Credit and Banking* in press.

Brown, C., J. Chiu, and T. V. Koepll (2019). “What drives Bitcoin fees? Using SegWit to assess Bitcoin’s long-run sustainability”. *Queen’s Economics Department working paper* 1423.

Budish, E. (2018). “The economic limits of bitcoin and the blockchain”. *National Bureau of Economic Research working paper* 24717.

Easley, D., M. O’Hara, and S. Basu (2019). “From mining to markets: the evolution of bitcoin transaction fees”. *Journal of Financial Economics* 134.1, pp. 91–109.

Ersoy, O., S. Roos, and Z. Erkin (2019). “How to profit from payments channels”. URL: <https://arxiv.org/pdf/1911.08803.pdf>.

Fruchterman, T. M. J. and E. M. Reingold (1991). “Graph drawing by force-directed placement”. *Software – Practice and Experience* 21.11, 1129–1164.

Hautsch, N., C. Scheuch, and S. Voigt (2018). “Limits to arbitrage in markets with stochastic settlement latency”. *Center for Financial Studies Working Paper* 616.

Huberman, G., J. Leshno, and C. Moallen (2017). “Monopoly without a monopolist: an economic analysis of the Bitcoin payment system”. *CEPR discussion paper* 12322.

Hyndman, R. J. and Y. Khandakar (2008). “Automatic time series forecasting: the forecast package for R”. *Journal of Statistical Software* 27.1, pp. 1–22.

Levine, J. (2019). “Scalability controversy: understanding past cryptocurrency returns through Segregated Witness”. *UC Santa Barbara honors thesis*.

Libra Association (2019). “An introduction to Libra”.

Makarov, I. and A. Schoar (2019). “Trading and arbitrage in cryptocurrency markets”. *Journal of Financial Economics* in press.

Nakamoto, S. (2008). “Bitcoin: a peer-to-peer electronic cash system”. *Mimeo*.

Pérez-Solà, C., S. Delgado-Segura, J. Herrera-Joancomartí, and G. Navarro-Arribas (2019). “Analysis of the SegWit adoption in Bitcoin”. *Mimeo*. URL: <http://www.deic.uab.es/~guille/publications/papers/2018.recsi.segwit.pdf>.

Poon, J. and T. Dryja (2016). “The Bitcoin Lightning Network: scalable off-chain instant payments”. URL: <https://lightning.network/lightning-network-paper.pdf>.

Schilling, L. and H. Uhlig (2019). “Some simple Bitcoin economics”. *Journal of Monetary Economics* 106, pp. 16–26.

Thakor, A. V. (2019). “Fintech and banking: what do we know?” *Journal of Financial Intermediation* in press.

Zimmerman, P. (2019). “Blockchain structure and cryptocurrency prices”. *SSRN* 3478498.

Table 6: Definitions of variables. Data are from January 1, 2017 to September 5, 2019 unless otherwise indicated.

Variable	Definition
Mempool txn count	Total number of unconfirmed transactions in the Bitcoin (BTC) mempool. Source: Hoenicke.
Mempool txn fees (USD)	Total fees in USD of pending unconfirmed transactions in the Bitcoin mempool. Source: Hoenicke.
Low fee txns (%)	Percentage of transactions in the Bitcoin mempool offering fees less than 10 satoshis per virtual byte (sat/B). 100 million satoshis = 1 bitcoin. Source: Hoenicke.
BCH/BTC txns	Ratio of total Bitcoin Cash transactions to total Bitcoin transactions each day. Source: Coin Metrics. Data start Aug 1, 2017.
BCH/BTC fees	Ratio of total value of Bitcoin Cash transaction fees to Bitcoin transaction fees (both in USD) each day. Source: Coin Metrics. Data start Aug 1, 2017.
SegWit txns (%)	Average daily percentage of Bitcoin transactions per block that use Segregated Witness (SegWit). Source: Bitcoin Visuals. Data start Aug 23, 2017.
Fees by SegWit txns (%)	Percentage of total daily fees paid by SegWit transactions. Source: txstats.com. Data start Aug 23, 2017.
Lightning Network channels	Number of active channels on the Lightning Network. Source: hashxp.org. Data from Jan 1, 2018.
Lightning Network capacity (USD)	Total value of active channels on the Lightning Network (in USD). Source: hashxp.org. Data from Jan 1, 2018.
30-day volatility	Rolling standard deviation of bitcoin returns from past 30 trading days. Source: Coin Metrics.
1-day price change	Rolling difference in log Bitcoin price between days $t - 1$ and $t - 2$. Source: Coin Metrics.
Mining intensity	Expected rate of block creation, measured as total miners' hash rate supplied by miners divided by average difficulty. Source: Coin Metrics.