Finance Research Letters 52 (2023) 103480

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

The Lightning Network: Turning Bitcoin into money ☆

Anantha Divakaruni [a], Peter Zimmerman [b,1]

[a] University of Bergen, Norway
[b] Federal Reserve Bank of Cleveland, United States of America

ARTICLE INFO

JEL classification:
E42
G10

Keywords:
Bitcoin
Blockchain
Cryptocurrency
Lightning Network
Payments

ABSTRACT

The Lightning Network (LN) is a means of netting Bitcoin payments outside the blockchain. We find a significant association between LN adoption and reduced blockchain congestion, suggesting that the LN has helped improve the efficiency of Bitcoin as a means of payment. This improvement cannot be explained by other factors, such as changes in demand or the adoption of SegWit. We find mixed evidence on whether increased centralization in the Lightning Network has improved its efficiency. Our findings have implications for the future of cryptocurrencies as a means of payment and their environmental footprint.

1. Introduction

Bitcoin was originally designed to serve as a "peer-to-peer electronic cash system" — that is, a reliable means of payment outside the control of centralized monetary authorities (Nakamoto, 2008). Since its introduction in 2009, Bitcoin has grown immensely in value, but still sees relatively little use as a means of payment (Makarov and Schoar, 2021). One important reason is that Bitcoin's blockchain technology imposes capacity constraints on processing transactions. These constraints allow Bitcoin to handle, on average, seven transactions per second, which compares unfavorably to centralized publicly available payment systems such as Visa or Mastercard.[1] When transaction demand is high, the processing limits mean that Bitcoin transactions can take a long time to settle. In recent years, many solutions have been proposed to resolve this so-called *scalability problem*, to help Bitcoin achieve its potential as a large-scale payments system.

One such solution is the *Lightning Network* (LN), which allows Bitcoin users to make payments outside the blockchain. Rather than inscribe every individual payment onto the blockchain, two individuals can open an LN channel and make bilateral payments. Once they have completed their payments, they can close the channel and settle the net amount. In principle, doing this requires only two transactions on the blockchain – one to open the channel, and another to close it – regardless of the amount settled or the number of underlying payments. In this way, adoption of the LN promises to reduce demand for blockchain space and ease congestion.

We find that adoption of the Lightning Network has led to a reduction in Bitcoin blockchain congestion and lower mining fees. The results are significant, both statistically and economically, and cannot be explained by changes in demand for blockchain space,

☆ The views stated herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System. This paper was previously circulated under the title "Ride the Lightning: Turning Bitcoin into Money". The authors wish to thank Huzaizhi Chen, Rodney Garratt, Engin Iyidogan, Carlos León, and Ahmed Sewaid, along with seminar audiences at the Bank of England, Boca Corporate Finance and Governance Conference, CEBRA, CEMLA, Paris Fintech & Crypto Seminar, Southern Finance Association, Southwestern Finance Association, and UCL P2PFISY 2020. Any errors are the responsibility of the authors.

E-mail addresses: anantha.divakaruni@uib.no (A. Divakaruni), peter.zimmerman@clev.frb.org (P. Zimmerman).

[1] Visa claims to able to handle over 24,000 transactions per second. See https://usa.visa.com/run-your-business/small-business-tools/retail.html.

https://doi.org/10.1016/j.frl.2022.103480
Received 7 June 2022; Received in revised form 6 October 2022; Accepted 6 November 2022
Available online 19 November 2022
1544-6123/

nor by other technological developments. We find limited evidence that greater centralization of the network is associated with lower fees. Our results suggest that the Lightning Network can help Bitcoin achieve greater scalability, allowing it to operate better as a payments system. According to our results, if the LN had existed in 2017, congestion could have been 84 percent lower.

Our analysis covers the period 1 January 2017 to 5 September 2019. Data limitations prevent us from extending our data set. The Lightning Network has continued to grow, doubling in size over 2021.[2] There has been institutional adoption, too. For example, Twitter allows tipping using the LN, among other payment methods.[3] El Salvador enables Bitcoin payments among its citizens using the Chivo Wallet, which features LN functionality (Alvarez et al., 2022). And several cryptocurrency exchanges have introduced support for the LN.[4] But recent episodes of high congestion, especially in early 2021, suggest that the LN is not yet a panacea. Moreover, settlement capacity is not the only reason why Bitcoin is not widely used as a means of payment (Yermack, 2015).

The development of the Lightning Network may have consequences for welfare. First, as Bitcoin becomes a more efficient payments system, users are better off. Their transactions settle more quickly and cheaply. Second, since fewer transactions need to be recorded on the blockchain, less memory and energy are needed to run a Bitcoin node. This saving lowers the cost of maintaining the blockchain, allowing more nodes to participate and making the system more secure against a double-spending attack. Third, by reducing fees, the LN reduces the incentive for Bitcoin miners to use large amounts of computing power, meaning less energy use and positive consequences for the environment.[5] Fourth, less blockchain congestion may mean lower barriers to arbitrage across cryptocurrency exchanges, thereby improving market liquidity (see Hautsch et al., 2018).

While this paper focuses on Bitcoin, all cryptocurrencies that use permissionless verification have transaction capacity constraints. Decentralized consensus is required to confirm new transactions, and it takes time for nodes around the world to communicate consensus to one another (Hinzen et al., 2022). Netting solutions like the Lightning Network can allow other cryptocurrencies to be widely used, secure, and decentralized. For example, the Raiden Network is a similar innovation for Ethereum. Other solutions to the scalability problem have been proposed, including sharding, and batching at exchange level. If the scalability problem can be successfully addressed, it may be possible for a currency based on a permissionless blockchain to obtain wide acceptance.

The rest of the paper is organized as follows. Section 2 briefly describes the Lightning Network and outlines findings from the existing literature. Section 3 describes our data, and Section 4 our results. Section 5 concludes.

2. The Lightning Network

The Lightning Network was first introduced by Poon and Dryja (2016), and began to attract widespread usage in January 2018. The LN is a secondary transaction layer that operates outside the blockchain. Two users open an LN channel by contributing Bitcoin to a smart contract. They can then transfer these coins between them without creating traffic on the blockchain (Auer, 2019). Once the channel is closed, only the net amount needs to be settled on-chain as a single payment. This netting reduces the required number of on-chain transactions to just two: one to introduce the smart contract that opens an LN channel, and a second to close it. In this way, the system can handle a much larger number of payments. Arcane Research (2022) provides an up-to-date description of the Lightning Network.[6]

The Lightning Network protocol itself relies on *Segregated Witness* (SegWit), a change to the Bitcoin transaction format that improves the efficiency of blockchain storage, so that a single Bitcoin block can potentially store up to four times as many transactions as before. Brown et al. (2021) show that the introduction of SegWit has reduced Bitcoin mining fees.

Only a couple of papers in the economics and finance literature focus on the Lightning Network: see Bertucci (2020) and Guasoni et al. (2021). More broadly, our paper relates to a literature examining the fee-based market for blockchain space; see, for example, Easley et al. (2019), Huberman et al. (2021), Lehar and Parlour (2020) and Makarov and Schoar (2020).

3. Data

We aim to test whether adoption of the Lightning Network is associated with reduced congestion on the Bitcoin blockchain. We construct measures of congestion using data on the Bitcoin *mempool*; that is, the set of payments waiting to be added to the blockchain. Our data come from Jochen Hoenicke.[7] We collect data on: (i) the number of pending transactions (*mempool txn count*); (ii) the fees attached to pending transactions (*mempool txn fees*); and (iii) the proportion of transactions with fees under 10 satoshis per virtual byte (*low fee txns*).[8]

---

2 See https://bitcoinvisuals.com/lightning.

3 See https://blog.twitter.com/en_us/topics/product/2021/bringing-tips-to-everyone.

4 See https://github.com/theDavidCoen/LightningExchanges.

5 The total energy consumption of Bitcoin miners is substantive, so the benefits could potentially be large. See https://ccaf.io/cbeci/index. On the other hand, a lower hash rate could make Bitcoin more vulnerable to double-spending attacks. In any case, these benefits and costs may not be realized immediately, because fees currently comprise a small part of miners' revenue and are expected to grow in importance over time (Easley et al., 2019).

6 Use of the LN can introduce new vulnerabilities. For example, users must be online while they are sending and receiving payments, putting them at risk of attack.

7 See https://jochen-hoenicke.de/queue. Hoenicke operates a Bitcoin node with its own mempool. There is no definitive mempool: each Bitcoin node may detect different pending payments. We assume that Hoenicke's data set is representative of all transactions pending in the Bitcoin network.

8 A *virtual byte* is equivalent to a physical byte for non-SegWit transactions and to four physical bytes for SegWit transactions. Since SegWit allows data to be stored up to four times as efficiently, a virtual byte is a measure of the amount of data encoded to the blockchain. Hoenicke only provides fees per virtual byte, not per physical byte. In addition, the data do not include transactions with zero fees. This omission is because it is costless for a vexatious attacker to submit zero-fee transactions to the mempool, so miners tend to ignore them. Including zero-fee transactions could therefore overstate the actual level of mempool congestion. Easley et al. (2019) study the determinants of zero-fee transactions in Bitcoin.

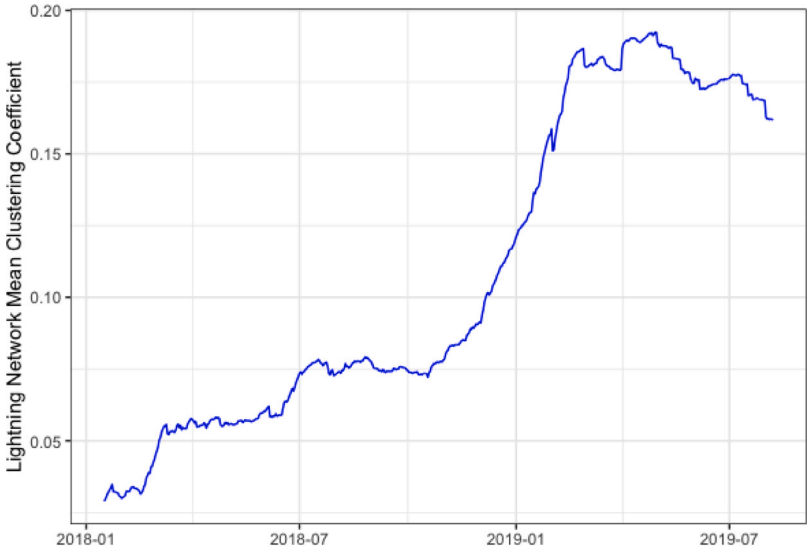

Fig. 1. Mean clustering coefficient among Lightning Network nodes.
Source: hashXP.

Data on the Lightning Network come from the website hashXP. 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, hashXP provides complete details of Bitcoin transactions executed in order to open and close LN channels.[9]

The shape of the Lightning Network may affect its efficiency. For example, if Lightning channels tend to be intermediated via a few central nodes, then collateral (i.e., the Bitcoin that users have locked into the LN) can be used more efficiently. In other words, when the network is more centralized, each channel, and each Bitcoin locked into the protocol, is likely to support a higher volume of payments (see Martinazzi and Flori, 2020). To account for this effect, we include the LN *clustering coefficient* as an independent variable. This network statistic is defined by Watts and Strogatz (1998) as the average probability that two neighbors of any given node are themselves connected. When the network is more centralized, the clustering coefficient is lower. Thus, we predict that when the clustering coefficient is high, mempool congestion is worse. As Fig. 1 shows, the network has tended to become more clustered – and thus less centralized – as it develops, though the last few months of the sample period show a trend toward greater centralization.

We introduce proxies for Bitcoin demand. Higher demand for transactions on the Bitcoin blockchain can increase congestion for reasons unrelated to LN adoption. Liu and Tsyvinski (2020) suggest that demand is positively related to historical price changes; in other words, there is a momentum factor. Motivated by this observation, we introduce both the daily closing Bitcoin price and the 1-day change in log-price. Prices are taken at midnight UTC each day. We also use price volatility as a proxy for speculative demand for Bitcoin, defining *30-day volatility* as the rolling standard deviation of Bitcoin returns from the past 30 trading days. These three measures are computed using price data from Coin Metrics (https://coinmetrics.io/).[10]

We also include a measure for the supply of blockchain space. Unlike demand, supply is directly observable ex post, since we can see how many blocks were created each day. We proxy supply by dividing miners' total hash rate divided by average mining difficulty; we call this measure *mining intensity*. While the Bitcoin protocol aims for a long-run mean of one block every 10 minutes, the realized rate of block creation can vary due to chance, or due to changes in miners' hash rate since the previous difficulty adjustment (Nakamoto, 2008).

Since SegWit adoption may affect mempool congestion, we control for it in our regressions. We obtain data from Bitcoin Visuals on the estimated proportion of Bitcoin transactions that use SegWit (https://bitcoinvisuals.com/chain-tx-block). A description of each variable can be found in the Appendix.

Our sample period contains daily data from 1 January 2017 to 5 September 2019, so it includes a period of about a year before the LN was widely adopted. We cannot extend our data set any later because, beyond these dates, hashXP was no longer actively

[9] See https://hashxp.org/lightning. To our knowledge, these data are available for public access and there is no restriction on their use. The data set contains only public LN channels. Users can also open private channels, which are known only to the connecting nodes and not announced to the broader network, but no public data are available on these private channels.

[10] We use a mixture of absolute price and logged price to minimize redundancy in the ARIMA structure. We use 30-day volatility as it is the most commonly used measure (e.g., Liu and Tsyvinski, 2020). As a robustness check, we find that using 24-hour volatility makes little difference to our results. We thank an anonymous referee for these suggestions.

Table 1
Summary statistics.

	count	mean	std dev	min	median	max
Mempool txn count	968	23,042	40,619	92	5731	252,750
Mempool txn fees (USD)	968	106,180	440,206	39	3008	4,750,619
Low fee txns (%)	968	53.45	28.30	0	52.04	95.99
Lightning Network channels	968	12,671	15,374	0	7575	44,087
Lightning Network capacity (USD)	968	2,766,535	4,080,066	0	205,388	11,794,337
Lightning Network mean clustering	968	0.06	0.07	0	0.06	0.19
SegWit txns (%)	968	20.72	15.48	0	27.61	46.80
30-day volatility	968	4.16	1.54	1.10	4.03	8.07
Daily price (close)	967	6077.31	3642.24	777.76	6209.00	19,497.40
1-day log-price change	967	0.00	0.04	−0.21	0.00	0.23
Mining intensity	968	7.49	0.82	3.98	7.51	9.79

Notes: Daily data from 1 January 2017 to 5 September 2019. See the Appendix for variable definitions and data sources.

monitoring the Lightning Network and providing accurate data. As a result, we are unable to study more recent developments in the LN.

Hoenicke's mempool data set is missing six days: 1 Feb 2017; 17–19 Apr 2017; 1 Jun 2019; and 26 Jun 2019. We drop these days from our data set. We use first-differenced data (as explained later in this section), so we also drop the following days (i.e., 2 Feb 2017; 20 Apr 2017; 2 Jun 2019; 27 Jun 2019). As a result, we have a total of 968 daily observations of the dependent variables. In addition, the Coin Metrics data on prices are missing one day (1 Jan 2019).

Table 1 shows summary statistics for our data. Many of the variables are highly volatile with substantive right-skew. Because of this skewness, we use the logarithms of mempool txn count, mempool txn fees, LN channels, and LN capacity in our regressions.

Fig. 2(a) plots the number of transactions waiting to be confirmed in Bitcoin's mempool (denoting congestion) over our sample period, along with active LN channels over time and the percentage of transactions that use SegWit. Congestion in Bitcoin has fallen markedly since early 2018, coinciding with the introduction and rapid adoption of the LN. Congestion has remained relatively low since then, although it picked up slightly in mid-2019.[11] Fig. 2(b) plots similar measures weighted by monetary value: we measure congestion using mempool fees, LN adoption using the USD value of locked Bitcoin, and SegWit usage by the monetary value of transactions. Total fees attached to payments waiting in the mempool have fallen since 2017, suggesting either lower demand or greater supply of settlement capacity. Over this period, the total value of Bitcoin used to collateralize LN channels has risen.

Fig. 3 shows that the distribution of fees has changed over our sample period. Generally, fees have fallen in nominal bitcoin terms. The proportion of transactions with fees below 10 satoshis per virtual byte rose from 32.6 percent on 1 January 2018 to 74.2 percent on 5 September 2019.[12]

We are interested in whether LN adoption is associated with lower mempool congestion. We test for relationships using autoregressive integrated moving average (ARIMA) specifications, which estimate the following regression equation:

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

where c is a constant term, y_t^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, d is the number of differences taken, and q is the length of the moving average window of historical residual terms. For each specification, we estimate the parameters (p, d, q) using the Hyndman–Khandakar algorithm (Hyndman and Khandakar, 2008). The time variable t is daily. We employ robust standard errors, since we cannot be sure of homoskedasticity.

Figs. 2(a) and 2(b) suggest that the data are non-stationary. We take first-differences of all our variables. Augmented Dickey–Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests confirm that these first-differenced variables are stationary (i.e., $d = 1$).

4. Results

We run three sets of regressions. First, we test the effect of LN adoption on mempool count, using the number of LN channels. We run four versions of this model. Model (1) contains no controls; model (2) includes the demand and supply controls; model (3) includes the proportion of transactions that use SegWit; and model (4) includes all the independent variables. Table 2 reports the

[11] By September 2019, the average daily mempool count was 75 percent lower than at the start of 2017. This decline in congestion does not appear to be driven by lower demand for Bitcoin. Although demand initially declined following the collapse of the cryptocurrency market in early 2018, the number of confirmed transactions subsequently grew to over 300,000 transactions per day by 5 September 2019, nearly back to its 2017 peak. See https://www.blockchain.com/charts/n-transactions.

[12] There are 100 million satoshis to a bitcoin.

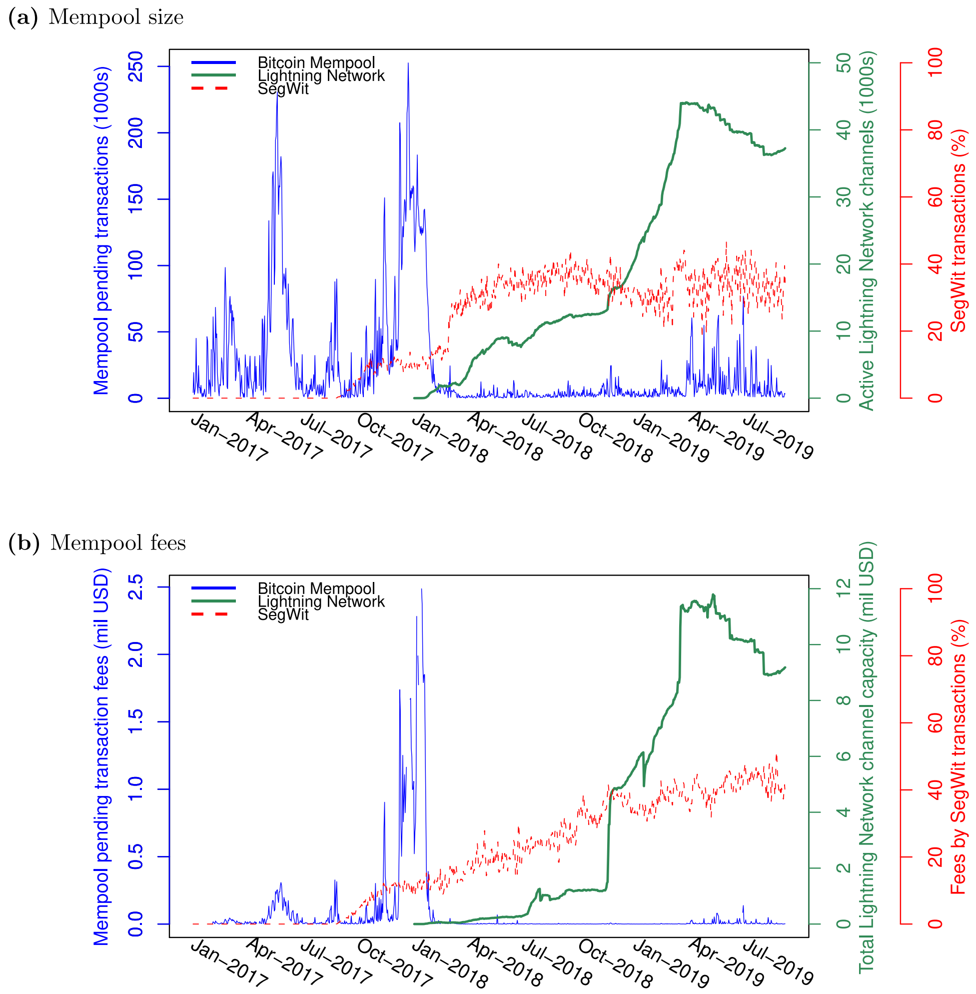

Fig. 2. Bitcoin mempool and the adoption of the Lightning Network and SegWit.*Notes*: Daily data from 1 January 2017 to 5 September 2019. *Source:* Jochen Hoenicke, blockchain.com, and hashXP.

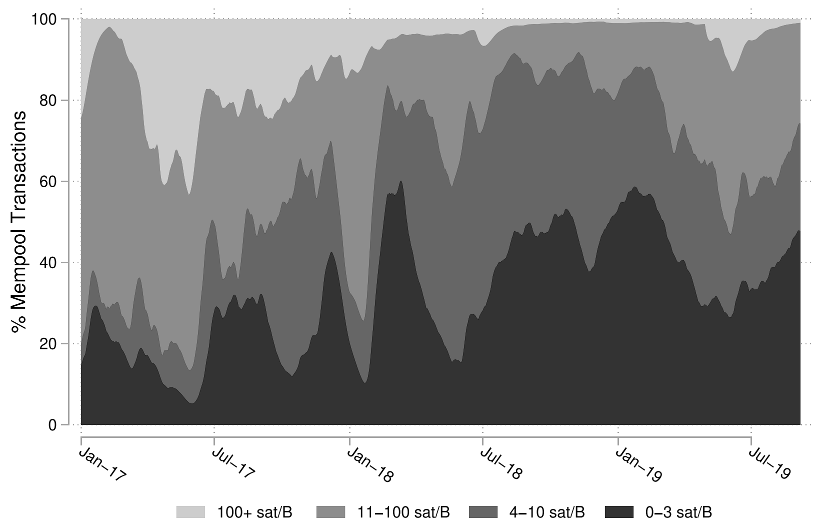

Fig. 3. Distribution of fees in the Bitcoin mempool.
Notes: Daily data from 1 January 2017 to 5 September 2019. The chart plots fees in satoshis per virtual byte. There are 100 million satoshis to a bitcoin.
Source: Jochen Hoenicke.

Table 2
Impact of Lightning Network adoption by number of channels on mempool count.

	(1)	(2)	(3)	(4)
ΔLN channels (log)	−0.247***	−0.168**	−0.251***	−0.174**
	(0.075)	(0.083)	(0.077)	(0.084)
ΔLN mean clustering	−3.857	−3.317	−4.004	−3.530
	(5.959)	(5.907)	(5.800)	(5.769)
ΔSegWit txns (%)			0.016	0.017
			(0.012)	(0.012)
Δ30-day volatility		−0.013		−0.009
		(0.077)		(0.078)
ΔDaily price		1.275***		1.249***
		(0.472)		(0.472)
Δ1-day log-price change		−1.383**		−1.376**
		(0.689)		(0.687)
ΔMining intensity		0.024		0.034
		(0.049)		(0.049)
Constant	−0.001	−0.005	−0.001	−0.005
	(0.009)	(0.009)	(0.009)	(0.009)
Observations	967	966	967	966
AIC	2589	2585	2589	2585
$p(Q)$	0.957	0.909	0.963	0.923

Notes: Regressions of LN channels (log) on mempool transaction count (log). In all four models, the Hyndman–Khandakar parameters are $p = 6, d = 1, q = 2$. Data are from 1 January 2017 to 5 September 2019. Standard errors are given in parentheses.

Table 3
Impact of Lightning Network adoption by capacity value on mempool fees.

	(1)	(2)	(3)	(4)
ΔLN capacity (USD log)	−0.198**	−0.088	−0.197**	−0.089
	(0.099)	(0.088)	(0.100)	(0.089)
ΔLN mean clustering	−6.889	−5.094	−7.150	−5.449
	(7.940)	(7.348)	(7.652)	(7.143)
ΔSegWit txns (%)			0.025*	0.024*
			(0.014)	(0.014)
Δ30-day volatility		0.120		0.129
		(0.090)		(0.090)
ΔDaily price		3.034***		3.008***
		(0.557)		(0.561)
Δ1-day log-price change		−2.103**		−2.089**
		(0.834)		(0.833)
ΔMining intensity		0.012		0.022
		(0.059)		(0.059)
Constant	0.003	−0.006	0.002	−0.007
	(0.010)	(0.008)	(0.010)	(0.008)
Observations	967	966	967	966
AIC	3042	3020	3040	3019
$p(Q)$	0.907	0.932	0.916	0.934

Notes: Regressions of LN capacity (USD log) on mempool fees (USD log). In all four models, the Hyndman–Khandakar parameters are $p = 6, d = 1, q = 1$. Data are from 1 January 2017 to 5 September 2019. Standard errors are given in parentheses.

results. In each of the four specifications, an increase in the number of LN channels reduces the mempool count. The results are significant at the 5 percent level. Higher Bitcoin prices tend to be associated with worse congestion, as we might expect.[13]

Our second set of results tests a similar relationship using US dollar-weighted values. We regress the USD value of Bitcoin locked into the LN against the USD value of fees attached to mempool transactions. Table 3 shows the results. As before, greater LN capacity is associated with reduced congestion. This time, however, the results are not significant once we include changes in prices. This suggests that changes in the dollar value of congestion can be well explained by changes in the day-to-day changes in the dollar value of Bitcoin.

Finally, we investigate how the LN affects the proportion of low fee transactions in the mempool. Table 4 shows that greater LN usage is associated with a significant increase in low fee transactions. Unlike the first two sets of regressions, clustering has a

[13] For each model, Ljung–Box and Durbin–Watson tests suggest no evidence of autocorrelation in the residuals. We record $p(Q)$, the p-value from the Ljung–Box test.

Table 4
Impact of Lightning Network adoption by number of channels on low fee mempool transactions.

	(1)	(2)	(3)	(4)
ΔLN channels (log)	0.192***	0.172***	0.195***	0.175***
	(0.034)	(0.038)	(0.034)	(0.037)
ΔLN mean clustering	−2.244***	−2.322***	−2.169**	−2.173**
	(0.858)	(0.867)	(0.864)	(0.875)
ΔSegWit txns (%)			−0.016***	−0.015***
			(0.004)	(0.004)
Δ30-day volatility		−0.018		−0.024
		(0.024)		(0.024)
ΔDaily price		−0.150		−0.140
		(0.154)		(0.152)
Δ1-day log-price change		0.135		0.136
		(0.233)		(0.231)
ΔMining intensity		0.043***		0.036**
		(0.015)		(0.015)
Constant	−0.002	−0.002	−0.002	−0.001
	(0.002)	(0.002)	(0.002)	(0.002)
Observations	967	966	967	966
AIC	692	675	676	660
$p(Q)$	0.984	0.995	0.997	1.000

Notes: Regressions of LN channels (log) on proportion of mempool transactions with fees below 10 satoshis per virtual byte. In all four models, the Hyndman–Khandakar parameters are $p = 6, d = 1, q = 2$. Data are from 1 January 2017 to 5 September 2019. Standard errors are given in parentheses.

significant and negative impact on low fees. In other words, a more centralized network means that transactions are likelier to have low fees, in line with our priors.

Overall, these results suggest that increased LN usage is associated with a significant reduction in mempool congestion. Since there is no theoretical upper limit on LN usage, there is the potential for further reductions in congestion in the future. However, network centralization does not have a clear effect on the efficiency of the network.

In each regression, SegWit has the opposite effect of Lightning Network adoption on mempool congestion, although the results are only significant in Table 4. At first glance, the signs of the coefficients are surprising: greater use of SegWit appears to increase, rather than reduce, congestion. There are a number of possible explanations. First, LN transactions require SegWit, so there is some positive correlation between these variables. However, the exact relationship is not clear, since we do not have data on the number of LN transactions, only on the number of channels and the value of Bitcoin locked in. Second, the causality is not clear. It may be that periods of high congestion incentivize greater SegWit usage. Third, since SegWit transactions use fewer virtual bytes than non-SegWit transactions (all else equal), users may be willing to pay a higher fee per virtual byte.[14]

We can assess the economic significance of reducing Bitcoin congestion by posing the following counterfactual question: if, during 2017, the LN had existed and been the size it was at the end of our sample, by how much would Bitcoin congestion have been lowered? Our results suggest that the mempool count would have been 84 percent lower, mempool fees 76 percent lower, and the proportion of low fee transactions 184 percent higher. These numbers demonstrate that the LN can potentially have a substantial impact on blockchain congestion.

5. Conclusion

We show that usage of the Lightning Network is associated with reduced mempool congestion in Bitcoin and with lower fees. Our findings suggest that the off-chain netting benefits of the Lightning Network can help Bitcoin to scale and function better as a means of payment. Centralization of the Lightning Network does not appear to make it much more efficient, though it may increase the proportion of low fee transactions.

As the Lightning Network expands, the potential grows for Bitcoin to operate better as a payment system, although we cannot say for sure whether this is happening. Data limitations prevent us from extending our analysis beyond 2019. And there are still episodes of high congestion, such as in early 2021, suggesting that the Lightning Network has not yet solved Bitcoin's scalability problem.

Data availability

Data will be made available on request.

[14] Suppose, for example, that Bitcoin users have the power to set fees, while miners are price takers. Then fees are unlikely to be very sensitive to the amount of blockchain space a transaction uses. In that case, a SegWit transaction may attract a similar fee to its non-SegWit counterpart, and so will appear to be more expensive per virtual byte.

Appendix. Definitions of variables

Variable	Definition
Mempool txn count	Total number of unconfirmed transactions in the Bitcoin (BTC) mempool. Source: Jochen Hoenicke.
Mempool txn fees (USD)	Total fees in USD of pending unconfirmed transactions in the Bitcoin mempool. Source: Jochen Hoenicke.
Low fee txns (%)	Percentage of transactions in the Bitcoin mempool offering a fee lower than 10 satoshis per virtual byte. 100 million satoshis = 1 bitcoin. Source: Jochen Hoenicke.
Lightning Network channels	Number of active channels on the Lightning Network. Data from 1 Jan 2018. Source: hashXP.
Lightning Network capacity (USD)	Total value of active channels on the Lightning Network (in USD). Data from 1 Jan 2018. Source: hashXP.
Lightning Network mean clustering	Mean clustering coefficient across Lightning nodes, as defined by Watts and Strogatz (1998). Source: hashXP.
SegWit txns (%)	Average daily percentage of Bitcoin transactions per block that use Segregated Witness (SegWit). Data from 23 Aug 2017. Source: Bitcoin Visuals.
30-day volatility	Rolling standard deviation of Bitcoin returns from past 30 trading days. Source: Coin Metrics.
Daily price	Closing Bitcoin price on day t. Source: Coin Metrics.
1-day log-price change	Rolling difference in log Bitcoin closing price between days $t-1$ and $t-2$. Source: Coin Metrics.
Mining intensity	Expected rate of block creation, measured as total hash rate supplied by miners divided by average difficulty. Source: Coin Metrics.
AIC	Akaike information criterion.
$p(Q)$	p-value from Ljung–Box test for serial autocorrelation.

Note: Data are from 1 January 2017 to 5 September 2019 unless otherwise indicated.

References

Alvarez, F.E., Argente, D., Patten, D.V., 2022. Are cryptocurrencies currencies? Bitcoin as legal tender in El Salvador. National Bureau of Economic Research working paper no. 29968 http://dx.doi.org/10.3386/w29968.

Arcane Research, 2022. The state of Lightning, volume 2. https://arcane.no/research/reports/the-state-of-lightning-volume-2.

Auer, R., 2019. Beyond the doomsday economics of proof-of-work in cryptocurrencies. BIS working paper no. 765.

Bertucci, L., 2020. Incentives on the Lightning Network: a blockchain-based payment network. Proceedings of December 2020 Finance Meeting EUROFIDAI–ESSEC http://dx.doi.org/10.2139/ssrn.3540581.

Brown, C., Chiu, J., Koeppl, T.V., 2021. What drives Bitcoin fees? Using SegWit to assess Bitcoin's long-run sustainability. J. Financ. Mark. Infrastruct. 9 (4), 1–25. http://dx.doi.org/10.21314/jfmi.2021.014.

Easley, D., O'Hara, M., Basu, S., 2019. From mining to markets: the evolution of bitcoin transaction fees. J. Financ. Econ. 134 (1), 91–109. http://dx.doi.org/10.1016/j.jfineco.2019.03.004.

Guasoni, P., Huberman, G., Shikhelman, C., 2021. Lightning Network economics: channels. Michael J. Brennan Irish Finance working paper no. 21–7 http://dx.doi.org/10.2139/ssrn.3840374.

Hautsch, N., Scheuch, C., Voigt, S., 2018. Building trust takes time: limits to arbitrage in blockchain-based markets. arXiv http://dx.doi.org/10.48550/arXiv.1812.00595.

Hinzen, F.J., John, K., Saleh, F., 2022. Bitcoin's limited adoption problem. J. Financ. Econ. 144 (2), 347–369. http://dx.doi.org/10.1016/j.jfineco.2022.01.003.

Huberman, G., Leshno, J., Moallemi, C., 2021. Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. Rev. Econom. Stud. 88 (6), 3011–3040. http://dx.doi.org/10.1093/restud/rdab014.

Hyndman, R.J., Khandakar, Y., 2008. Automatic time series forecasting: the forecast package for R.. J. Stat. Softw. 27 (3), 1–22. http://dx.doi.org/10.18637/jss.v027.i03.

Lehar, A., Parlour, C.A., 2020. Miner collusion and the Bitcoin protocol. SSRN working paper no. 3559894 http://dx.doi.org/10.2139/ssrn.3559894.

Liu, Y., Tsyvinski, A., 2020. Risks and returns of cryptocurrency. Rev. Financ. Stud. 34 (6), 2689–2727. http://dx.doi.org/10.1093/rfs/hhaa113.

Makarov, I., Schoar, A., 2020. Trading and arbitrage in cryptocurrency markets.. J. Financ. Econ. 135 (2), 293–319. http://dx.doi.org/10.1016/j.jfineco.2019.07.001.

Makarov, I., Schoar, A., 2021. Blockchain analysis of the Bitcoin market. National Bureau of Economic Research working paper no. 29396 http://dx.doi.org/10.3386/w29396.

Martinazzi, S., Flori, A., 2020. The evolving topology of the Lightning Network: centralization, efficiency, robustness, synchronization, and anonymity. PLoS One 15 (1), http://dx.doi.org/10.1371/journal.pone.0225966.

Nakamoto, S., 2008. Bitcoin: a peer-to-peer electronic cash system. Mimeo. https://nakamotoinstitute.org/bitcoin/.

Poon, J., Dryja, T., 2016. The Bitcoin Lightning Network: scalable off-chain instant payments. https://lightning.network/lightning-network-paper.pdf.

Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of 'small-world' networks. Nature 393, 440—442. http://dx.doi.org/10.1038/30918.

Yermack, D., 2015. Is Bitcoin a real currency? An economic appraisal. In: Lee Kuo Chuen, D. (Ed.), Handbook of Digital Currency. Academic Press, San Diego, ISBN: 978-0-12-802117-0, pp. 31–43. http://dx.doi.org/10.1016/B978-0-12-802117-0.00002-3.