Lightning Network Economics: Channels*

Paolo Guasoni † Gur Huberman ‡ Clara Shikhelman §

January 9, 2023

Abstract

Compared with existing payment systems, Bitcoin's throughput is low. Designed to address Bitcoin's scalability challenge, the Lightning Network (LN) is a protocol allowing two parties to secure bitcoin payments and escrow holdings between them. In a lightning channel, each party commits collateral towards future payments to the counterparty and payments are cryptographically secured updates of collaterals. The network of channels increases transaction speed and reduces blockchain congestion. Focusing on a single channel in isolation from the LN, this paper (i) identifies conditions for two parties to optimally establish a channel, (ii) finds explicit formulas for channel costs, (iii) obtains the optimal collaterals and savings entailed, and (iv) derives the implied reduction in congestion of the blockchain. Unidirectional channels costs grow with the square-root of payment rates, while symmetric bidirectional channels with their cubic root. Asymmetric bidirectional channels are akin to unidirectional when payment rates are significantly different, otherwise to symmetric bidirectional.

*Partially supported by SFI (16/IA/4443, 16/SPP/3347), Columbia-IBM Center for Blockchain and Data Transparency, Chaire Fintech at University Paris Dauphine – PSL, The Algorand Foundation, and Simons Institute. We thank Jacob Leshno, Jiasun Li, Julien Prat, Fahad Saleh, and seminar participants at Harvard University, SIAM FME, Simons Institute, the Technion, and the 2022 CBER Conference for useful comments.

†Dublin City University, School of Mathematical Sciences, Glasnevin, Dublin 9, Ireland, email paolo.guasoni@dcu.ie

‡Columbia Business School, Kravis Hall, New York, NY 10027, USA, email gh16@columbia.edu

§Chaincode Labs 450 Lexington Avenue, New York, NY 10017, email clara.shikhelman@gmail.com

1 Introduction

To economize on transaction costs, parties to frequent transactions often arrange to pay for them periodically rather than immediately after each transaction. For instance, credit card holders are billed monthly. Similar arrangements are emerging for payments in blockchain-based cryptocurrencies such as Bitcoin and Ethereum. In these arrangements the parties deploy cryptographic tools to guarantee the payments. The guarantees are off-chain. Periodically, the parties settle their obligations on-chain. Consequently, (i) the parties economize on transaction costs, and (ii) the system's throughput improves thanks to the shift of interactions off-chain, thereby also reducing the time it takes until a transaction is practically irreversible. "Layer-two solutions" is a label for protocols which cryptographically secure payments off-chain and settle on-chain when necessary.

Examples of layer-two payment are readily available: sovereign-issued money backed by gold; commercial bank-issued money backed by deposits with the central bank; credit card-based payments backed by banks' payment networks; gold and silver deposit certificates in the 17th and 18th century Bank of Amsterdam, which were used to settle transactions (Frost et al., 2020). (Due to their convenience, most of the time these certificates traded at 5% premium to the underlying metal.) In fact, the architecture of stablecoins is based on a similar idea: The issuer maintains a fund of fiat currency which backs the stablecoin 1:1.

The protocol of the Bitcoin Payment System (Nakamoto, 2008) publishes all transactions on the blockchain. By design, the protocol can handle at most a few transactions per second, which is orders of magnitude lower than most credit-card payments systems, often resulting in significant delays and transaction fees. The Lightning Network (henceforth, LN) is a layer-two payment solution which addresses these weaknesses (see Poon and Dryja (2015) and Wirdum (2016)). It is a cryptographically secured protocol for escrow holdings of Bitcoin and changes in the holdings (i.e., payments). The protocol also specifies the circumstances in which the parties' balances settle on-chain.

Channels are the basic building blocks of the LN. Functionally, a channel is a jointly held Bitcoin account which opens with the two holders' balances reflecting their initial, on-chain contributions. Over time, the channel holders update the balances to reflect payments between them. Balance updating leaves the sum of the balances intact. Balance updating renders payments immediately irreversible. Payments can be routed through a chain of channels in the LN by updating balances accordingly. However, funds can be used to pay third parties that are not on the LN only after the transactions are recorded on-chain.

In general, where bitcoin is a common and frequently used medium of exchange, the LN is a good candidate to reduce the load on the bitcoin blockchain and increase the system's overall efficiency. In that case, the LN would be used for everyday transactions, such as grocery shopping, paying for transportation, and other routine payments. In such a world, a customer would open a single channel to a well-connected node, and will route payments over the LN for their everyday use.

Already today, the sending of remittances, especially cross-border remittances, appears to be a particularly cost-effective application of the LN. Moreover, using bitcoin for small amount remittances seems to be gaining traction in less developed economies (von Luckner et al. (2021), Ibaba et al. (2021), Tetek (2021)).

Another contemporary use case is e-sports, recreational and professional video game

playing. This growing field brings together people from all over the world to watch and take part in video game competitions. The winners of these competitions are rewarded handsomely for their achievements. To make these awards available to people that are not easily connected to the western banking system, startups build infrastructure over the LN to allow for the prizes to be sent over the LN.[1] Unlike in-game coins or other solutions, the prize sent over the LN can be used outside a specific game to shop online, and often can be exchanged to the local currency.

In the steady-state model studied here, as soon as a balance which started positive is exhausted, both holders record their balances on-chain, close the channel, and reopen another one immediately.[2] The initial balance of each party is the same in each reopening of the channel. If user pairs expect payments to flow only from one of them to the other but not in reverse, then the channel is called *unidirectional*. If the users expect payments to flow in both directions (possibly at different rates), the channel is called *bidirectional*.

Parties who use a channel for multiple transactions secure the transactions immediately off-chain. They burden the blockchain only to open and close the channel, thereby avoiding multiple fee payments associated with on-chain transactions. However, using the channel is costly because it requires locking up funds inside the channel, thereby foregoing alternative usage. Intuitively, using a channel is a good idea when transactions are sufficiently frequent.

The costs of LN channels and their implications for the trade-off between on-chain and the LN payment alternative are the focus of this paper. Particular attention is given to cases in which each of the two parties is both a payer and a payee and the traffic from one party to the other offsets the traffic in the opposite direction.

This paper's approach may also be useful to the analyses of other layer-two solutions for Bitcoin, for other cryptocurrencies such as Ethereum and perhaps other payment systems. A companion paper (Guasoni et al., 2021) discusses implications of this paper's analysis to the topology of the LN.

1.1 Model and Contribution

Two parties (or nodes) share a channel after they open an on-chain joint account funded by balances (or collaterals) the parties contribute. A payment of X units from node 1 (Alice) to node 2 (Bob) translates into a reduction of the balance of node Alice by X units, and a corresponding increase in the balance of node 2. For the sake of tractability, the paper's results are obtained assuming a unit transaction size, i.e., $X = 1$. This convenient simplification also offers a simple heuristic for the more general setting of random transaction sizes: if transaction sizes are IID with arrival rate λ and mean transaction size ν – independent of arrival times – then the formulas remain approximately valid up to replacing λ with $\lambda\nu$. Put differently, the payment rates λ in the paper should be thought of the product of number of transactions per unit of time times the average transaction size. When thinking about the channel as a LN component, λ represents the overall transaction flow through the channel in

[1] See, for example, https://zebedee.ioZebedee.

[2] There are other solutions to rebalancing channels, yet if there are no funds in the LN these solution include on-chain transactions and entail similar fees. After a channel's closure, one could try to save on committed capital by delay reopening until the next transaction. However, this choice would sacrifice the immediacy of the next transaction, therefore it is ruled out to ensure that all transactions are treated equally.

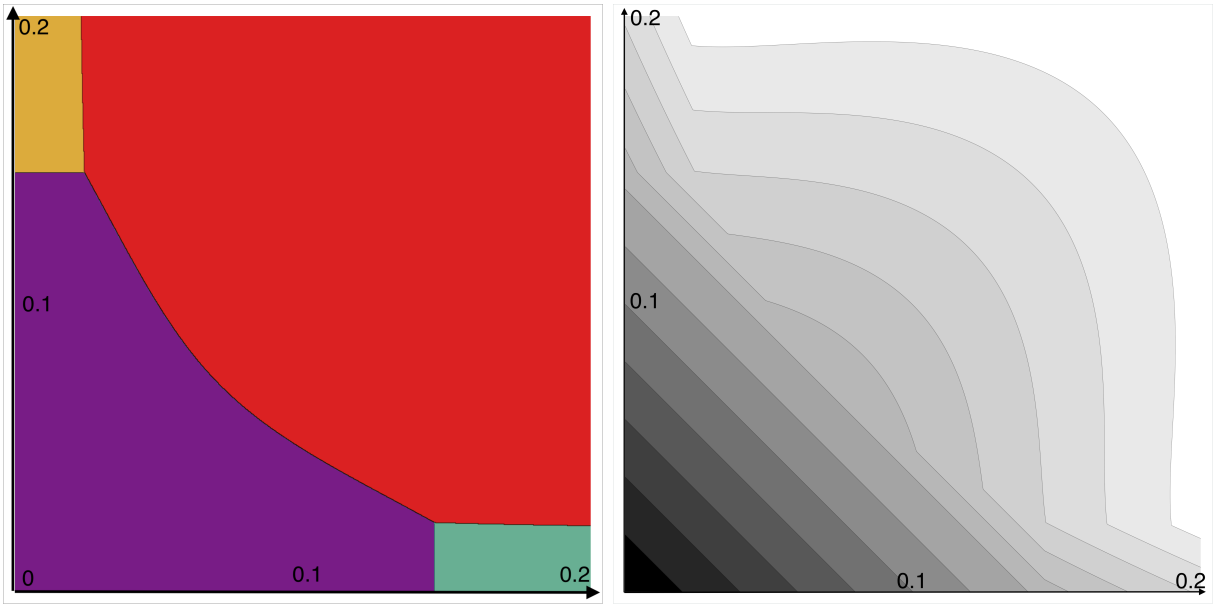

Figure 1: Optimal payment network (left) and its cost (right) if Alice pays Bob at rate λ_1 (horizontal axis, from 0 to 0.2), and Bob pays Alice at rate λ_2 (vertical axis, from 0 to 0.2 in number of annual transactions) with round-trip transaction cost $B = 1$ and interest rate $r = 1\%$. If both payment rates are small (purple, bottom left), all transactions optimally take place on chain, without lightning channels and the cost is proportional to the sum of such rates. If one rate is low but the other one is high (yellow and cyan, top left and bottom right), then the more frequent payer uses a unidirectional channel to pay the less frequent payer. The less frequent payer pays over the channel if his balance can support the payment and on-chain otherwise. If both rates are large enough (red, top right), then all payments take place through a bidirectional channel, and the cost is rather insensitive to an increase in the lower payment rate.

a given direction. That transaction flow aggregates the following: (i) flows originating from and destined to the two nodes of the channel; (ii) flows either originating from or destined to one of channel's nodes; (iii) flows originating from and destined to nodes outside the channel.

Demand for payments from node 1 to 2 (2 to 1) arrives at Poisson-governed rate λ_1 (λ_2); The continuously compounded discount rate is r. The cost of rebalancing a channel is B, and the cost of an on-chain transaction is C.[3] The analysis focuses on the channel cost assuming (i) the nodes choose to deposit initial balances of l_1 and l_2; (ii) the residual balances are posted on-chain; (iii) the parties renew the process after a node depletes its balance. The channel is unidirectional if λ_1=0 or $\lambda_2 = 0$; it is symmetric if $\lambda_1 = \lambda_2$.

This paper contributes to the Lightning Network literature by:

(i) offering a lean model of a channel in the LN. The model is close to but different from the well known Baumol (1952); Tobin (1956) and Miller and Orr (1966) models of demand for money and cash management;

[3]Huberman et al. (2021, 2019) argue that the transaction fee depends on user type and congestion level, which may vary with time. The present paper abstracts from these considerations for the sake of tractability and to focus on the tradeoff between transaction costs and capital opportunity costs.

(ii) solving for the channel cost as a function of the exogenously specified parameters and the parties' chosen initial balances (Theorem 3.2);

(iii) obtaining, in the realistic case of small interest rates, the cost minimizing initial balances (and the costs themselves) for unidirectional and symmetric channels.[4], extending the results for the unidirectional (respectively, bidirectional) channel to the nearly unidirectional (respectively, nearly symmetric bidirectional) case (Proposition 3.4);

(iv) showing that an asymmetric bidirectional channel is more akin to a unidirectional channel than to a symmetric channel (Proposition 3.5).

(v) establishing necessary lower bounds for the transaction frequencies to justify the existence of unidirectional and bidirectional channels (Theorem 3.3).

(vi) calculating the probability that one node exhausts its balance before the other (Proposition 3.7).

(vii) calculating the long-run ratio between the number of channel transactions and the number of on-chain transactions (Proposition 3.6).

A necessary step in the design and application of a channel is the comparison of a channel's cost with that of transacting on-chain. Such a comparison is particularly helpful if accompanied by the calculation of the cost-minimizing initial balances of the channel. This paper provides all these quantities.

The analysis supports and quantifies the initial intuition that a channel between two parties cuts on transaction costs if the transaction frequency is high enough. The benefit is highest when the transaction frequency in both directions is the same or almost the same. The benefit is more modest for unidirectional channels. Moreover, channels in which transaction frequency is not almost symmetric are akin to unidirectional channels.

Figure 1 illustrates some of the main findings: it displays the optimal (least costly) arrangement for two nodes, 1 and 2, who pay each other at rates λ_1 and λ_2. Five configurations can emerge: (i) All payments are on-chain; (ii) (respectively, (iii)) node 1 (resp. 2) pays on-chain when the channel cannot facilitate the transaction whereas the other node pays through a channel at all times; (iv) each node pays the other through a separate unidirectional channel; (v) both nodes use a single bidirectional channel. (By design, both panels are symmetric around the main diagonal.)

The left panel shows the four regions of the pairs (λ_1, λ_2). Relatively high transaction frequencies give rise to bidirectional channels. When transaction frequencies in both directions are low, no channel is used, as settlement occurs on-chain. When transaction frequencies are low in one direction and high in the opposite direction, the parties use a unidirectional channel to accommodate the latter, settling reverse transactions on-chain.

One message of this figure is that, for a channel-resetting cost equal to the size of each transaction, a symmetric bidirectional channel is the most economical choice, even for very modest frequencies, such as once every five years (i.e., $\lambda = 0.2$). A party that does not pay even as sporadically should pay on-chain while receiving through a unidirectional channel.

[4]Recall that the secular average of real short-term rates is less than 1%.

Only two parties that do not expect to transact in several years will shun lightning channels at all.

The right panel shows the equal-cost contours under the cost-minimizing behavior. The contours at the bottom left correspond to the lines $\lambda_1 + \lambda_2 = constant$, reflecting the linearity in volume of on-chain settlement. To interpret the other contours it is easiest to follow the uppermost contour, which is also the rightmost contour, from left to right. At the top there is a small flat portion corresponding to node 1 paying on-chain and node 2 paying through a unidirectional channel. Moving right, there is a small and sharp drop, indicating the transition from on-chain transactions to a hybrid of on-chain and unidirectional channel transactions. Moving right again – and this is the interesting part – the contour is approximately upward sloping for a while. In this part, payment frequency from node 1 to node 2 increases while reducing total cost, due to the benefits of netting. As λ_1 increases, the frequency of on-chain transactions comes down, thereby reducing the channel's total cost. This effect wears off as λ_1 comes closer to λ_2, hence the curved downward slope near $\lambda_1 = \lambda_2$.

1.2 Previous Work

Baumol (1952), Tobin (1956) (and, earlier but less well known, Allais (1947)) develop a model of transactions demand for money. The celebrated Baumol-Tobin work inspires the present paper's analysis of the unidirectional channel. Similarly, the seminal Miller and Orr (1966) model of a firm's demand for money inspires the present paper's analysis of the bidirectional channel. The inspiration of Baumol-Tobin and Miller-Orr notwithstanding, it is noteworthy that in these models one must transact in cash yet cash cannot be an exclusive store of value. In contrast, here there is always the possibility to execute transactions on-chain and use it to store value, thereby avoiding the LN and its channels.

Central to this paper is the tradeoff between the cost of capital and transaction fees. Huberman et al. (2021, 2019) propose and analyze a model in which the throughput of the Bitcoin Payment System is fixed by its protocol and transaction fees emerge as the system's response to congestion. Congestion requires the allocation of processing priority. This allocation is through transaction fees, which are higher when the system is more congested. An extension of the present work would be to the environment considered by Huberman et al. (2021, 2019).

A recent economic perspective of LN channels is offered by Brânzei et al. (2022). Assuming a cost function somewhat different from ours, they focus on symmetric bidirectional channels. The present paper develops theoretical results in the full generality of asymmetric channels of arbitrary size and transaction rates. It is also interesting to note another second-layer solution, similar to the LN, sketched in Narayanan et al. (2016).

2 An overview of the Lightning Network

The Lightning Network (LN) consists of channels and nodes. A channel has two participating nodes. A channel supports a series of bitcoin balance updates, i.e., payments between the channel's two participants. These updates are accomplished off the blockchain.

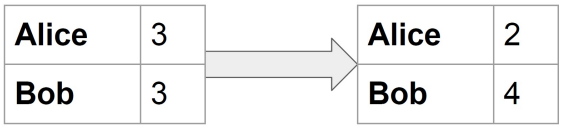

Figure 2: An example of Alice sending Bob a single bitcoin over their channel

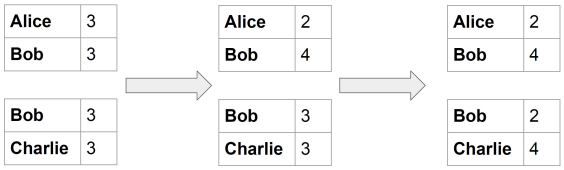

Figure 3: An example of Alice sending Charlie a single bitcoin through Bob

For example, Alice and Bob can commit 3 bitcoin each to a channel. When Alice wants to pay Bob 1 bitcoin she sends him a signed pending transaction that assigns her 2 bitcoin and assigns Bob 4 bitcoin (see figure 2). This transaction will not be sent to the blockchain.

Alice and Bob continue to update the balances. They will need to close and reopen the channel when Alice (or Bob) wants to pay an amount in excess of her (or his) current balance in the channel. As the cost of opening and closing a channel exceeds that of an on-chain transaction, a channel's goal is to support several transactions. A channel requires a commitment of funds, which cannot be used for other purposes while the channel is open.

The LN also supports payments between participants who do not have a channel in common but are linked through a path of channels. A node along the path may charge a fee for enabling the payment to go through the channels it participates in. For example, consider channels [A,B] between Alice and Bob and [B,C] between Bob and Charlie, such that each participant in each channel has a balance of three units. For Alice to send one unit to Charlie, she reduces her balance in the [A,B] channel by one unit in favor of Bob, who reduces his balance by a unit in the [B,C] channel in favor of Charlie. See figure 3 for an illustration.

In fact, a sequence of channel pairs can form a payment chain if each channel pair has a participant in common. For example channels [A,B], [B,C], [C,D] can jointly support payments between A and D. Cryptographic protocols guarantee that no party can be harmed by entering this arrangement. See Wirdum (2016) for an overview and Poon and Dryja (2015) for the full technical details.

In a normal course of events, for each channel there are only two on-chain transactions, namely, opening and closing the channel, i.e., committing and releasing funds, respectively. In addition, and off chain, participants update the balances as the need arises.

If one of the parties in a channel tries to cheat, the other can punish it by taking all of the funds in the channel. This is ensured by a cryptographic scheme detailed below.

2.1 Cryptographic Foundation

To open a lightning channel, each party: (i) deposits some bitcoin to a multiple signature (henceforth, multisig) address, co-signed by both and recorded in the blockchain; (ii) creates

a public-secret key pair and sends the public key to the other; (iii) creates a commitment transaction, recognizing the ownership of the respective amounts. For example, Alice deposits 12 bitcoin to the multisig address, while Bob deposits 8.[5] By co-signing, Alice and Bob can claim their respective balances through the blockchain. Then Alice creates a public-secret key pair A^*_{public}, A^*_{secret} and sends A^*_{public} to Bob. Likewise, Bob creates B^*_{public}, B^*_{secret} and sends B^*_{public} to Alice. Alice creates a new public-secret key pair A^0_{public}, A^0_{secret} while Bob creates B^0_{public}, B^0_{secret} and use these keys to set up their first commitment transaction. This transaction attributes the original bitcoin to themselves, as follows: Alice attributes 12 bitcoin to herself (the original amount she put in the channel), and the remaining 8 to a special multisig address. Then, she sends A^0_{public} to Bob. Alice's special multisig address works as follows: funds can be spent on-chain cooperatively or non-cooperatively. A cooperative closure of the channel entails a new transaction, which gives each party its respective share, and is published in the blockchain.

If Bob does not cooperate, Alice can send the transaction she created to the Blockchain. If she does this, Bob can claim his 8 bitcoin immediately by using the secret B^*_{secret}, while Alice needs to wait 1000 blocks before she can move her funds. If Alice does not cooperate, Bob can do the same. In general, the funds of the party that withdraws unilaterally are delivered with a delay, enabling the other party to withdraw the rest of the funds sooner. The LN channel between Alice and Bob is open once (i) the opening transaction is on the Blockchain, (ii) public keys A^*_{public} and B^*_{public} are exchanged, and (iii) the commitment transactions are complete. Alice sends 1 bitcoin to Bob over the LN as follows. First, Alice and Bob create new public-secret key pairs, associated with these specific balances. Alice creates A^1_{public}, A^1_{secret} and Bob creates B^1_{public}, B^1_{secret}, then they exchange the public keys.

Alice creates a new transaction that attributes 11 bitcoin to herself, and 9 to a special multisig address. She then sends the signature for this transaction to Bob. Bob creates a similar transaction that gives him 9 bitcoin and sends 11 bitcoin to a special multisig address, and sends it to Alice. Neither transaction is broadcast to the Blockchain.

The payment is final once Bob has the guarantee that Alice will not attempt to broadcast the previous transaction, in which he receives only 8 bitcoin. To provide such a guarantee, Alice gives Bob her previous secret key A^0_{secret}. Now, if Alice tries to cheat, claiming the previous balance on-chain, Bob can take both his 8 bitcoin and Alice's 12 bitcoin. Likewise, Bob also sends Alice his previous secret key B^0_{secret}. As Alice is the one benefiting from the old balance, she must send her secret key first. Alice and Bob can keep updating the state of the channel by creating new key pairs and transactions, sharing old secret keys. They do not need to trust each other because, at any point in time, if one of them tries to cheat, the other one can claim all the funds in the channel. Conversely, each party can unilaterally claim his current balance, with some delay.

3 The Cost of a Lightning Channel

Consider two nodes, 1 and 2, which exchange payments at different rates: node 1 sends one unit of currency to node 2 at rate λ_1, in that the cumulative number of payments from node

[5] Channel opening is transitioning from supporting only unilateral funding to supporting bilateral funding, which is discussed in this paper.

1 to node 2 by time t is described by a Poisson process $N_t^{\lambda_1}$ with rate λ_1. Likewise, the payments from node 2 to node 1 are described by another Poisson process $N_t^{\lambda_2}$. The two Poisson processes are independent. To settle these two streams of payments, consider the following payment possibilities, which can be implemented through a blockchain with a LN.[6]

The first and simplest option is to make all transactions on-chain, without using lightning channels. The advantage of this choice is that it does not require the commitment of any capital locked inside a channel. The disadvantage is that each payment incurs the fixed cost C of an on-chain transaction. Intuitively, such an arrangement may be optimal only if the payment rates are very low.

Second, each paying node could establish a unidirectional channel to settle each stream of payments. This arrangement is more attractive when payment rates are sufficiently high. In this case, costs are lower when payments are made through a unidirectional channel, in which payment commitments are made against the payer's outstanding balance. When the balance is exhausted, the payments settle and the balance is replenished on-chain. The channel size (i.e., the amount committed by the paying node) is chosen to minimize cost. The drawback of two unidirectional channels is that they forego any savings from offsetting payments, which can be substantial if both payment rates are large enough. As shown shortly, savings are higher as transaction rates approach each other. Having a unidirectional channel to support payments in one direction and making on-chain payments in the opposite direction when the channel cannot support them could be cost minimizing for highly asymmetrical payment rates.

Third, both nodes could agree to establish a bidirectional channel, with each of them possibly committing different amounts. This option is the most flexible, in that the contributions of each node can be optimized in relation to both incoming and outgoing rates of payment, and the number of on-chain transactions is reduced by offsetting payments. The disadvantage is that such savings may not materialize if at least one of the two payment rates is small enough.

The allocation of the costs to the two nodes is a separate issue, which this paper does not address. For the sake of concreteness, the presentation below assumes that each party contributes its committed balance.

3.1 Exact Costs

To examine quantitatively these tradeoffs, it is convenient to start by considering the cost of settling a stream of payments with rate λ simply through on-chain transactions, i.e., without any lightning channels. In addition to payment rates, the critical quantities necessary to perform the analysis are C, the cost of an on-chain transaction, B, the cost of resetting a channel, and r, the continuously compounded interest rate, which represents the opportunity cost per unit of time of using a unit of capital for another purpose (including another channel).[7]

[6]Some implementations of payment channels entail limits on the number of payments or the number of times that payment flows can switch. The present analysis abstracts from some limitations, assuming that a channel remains viable as long as the balance of both parties is above zero.

[7]In the bitcoin network, one can reset a channel through *two* on-chain transactions, but we allow for the existence of cheaper alternatives within the lightning network, by keeping the costs B and C independent.

Lemma 3.1 (On-chain Cost)**.** *The on-chain cost for a transaction stream with rate λ is $C\lambda/r$.*

The next step is to evaluate the cost of unidirectional and bidirectional channels.

Theorem 3.2 (Exact Channel Costs)**.** *Let λ_1 be the payment rate from node 1 to node 2 and λ_2 the payment rate from node 2 to node 1, and assume that $\lambda_1 \leq \lambda_2$. If node 1 commits an amount l_1 to the channel and node 2 commits an amount l_2, then:*

(i) A unidirectional channel costs

$$L^{0,l_2}(0,\lambda_2) = l_2 + B\left(\left(\frac{r+\lambda_2}{\lambda_2}\right)^{l_2} - 1\right)^{-1}. \tag{1}$$

Its minimal cost, setting $k = B\log(1+r/\lambda_2)$, is

$$L^{opt}(0,\lambda) = \frac{B}{2k}\left(-k + \sqrt{k(k+4)} + 2\log\left(\frac{1}{2}\left(k + \sqrt{k(k+4)} + 2\right)\right)\right) \tag{2}$$

and is achieved for $l_2 = \frac{B}{k}\log\left(\frac{1}{2}\left(k + \sqrt{k(k+4)} + 2\right)\right)$.

(ii) A bidirectional channel costs

$$L^{l_1,l_2}(\lambda_1,\lambda_2) = l_1 + l_2 - B\frac{\alpha_-^{l_1}(1-\alpha_+^{l_1+l_2}) - \alpha_+^{l_1}(1-\alpha_-^{l_1+l_2})}{\alpha_-^{l_1}(1-\alpha_+^{l_1+l_2}) - \alpha_+^{l_1}(1-\alpha_-^{l_1+l_2}) + \alpha_+^{l_1+l_2} - \alpha_-^{l_1+l_2}}, \tag{3}$$

where

$$\alpha_\pm = \frac{\lambda_1 + \lambda_2 + r \pm \sqrt{(\lambda_2-\lambda_1)^2 + r^2 + 2r(\lambda_2+\lambda_1)}}{2\lambda_2}. \tag{4}$$

Remark The unidirectional formula (1) follows from the bidirectional formula (3) by substituting $\lambda_1 = l_1 = 0$. However, it is convenient to consider it separately, in view of its different asymptotic properties, as explained below.

This theorem offers a closed-form expression for a channel's cost, given the nodes' commitments l_1, l_2. The optimal values of such commitments are not available explicitly for a bidirectional channel, hence more sophisticated arguments are required to understand the conditions under which different types of channels are optimal.[8]

The next result demonstrates that when payment rates are low enough, neither unidirectional not bidirectional channels should be used.

Theorem 3.3 (Bounds on Payment Rates)**.**

(i) If a unidirectional channel with rate λ costs less than on-chain transactions, then

$$C > \log\left(1 + B\log\left(1 + \frac{r}{\lambda}\right)\right) = \frac{Br}{\lambda} + o(r). \tag{5}$$

[8]Note that the expression for $\alpha_\pm$ is ostensibly asymmetric in λ_2, λ_1 but the expression for the cost $L^{\lambda_2,\lambda_1}(l_1,l_2)$ is in fact symmetric, that is, invariant to swapping (λ_1,l_1) and (λ_2,l_2).

(ii) If a bidirectional symmetric channel with rates λ costs less than on-chain transactions, then

$$C > 3\left(\frac{B\kappa^4}{4}\right)^{1/3} - \frac{B}{12}\kappa^2 = 3\left(\frac{Br^2}{4\lambda^2}\right)^{1/3} - \frac{Br}{12\lambda} + o(r) \tag{6}$$

where $\kappa = \log\frac{r+2\lambda+\sqrt{r(4\lambda+r)}}{2\lambda}$.

The main message of Theorem 3.3 reflects the intuition that both unidirectional and bidirectional channels are optimal only when on-chain costs are high, channel-reset costs are low, payment rates are high, or interest rates are low. For this reason, Theorem 3.3 supports the asymptotic analysis in the limit of r near zero, as this is the relevant regime for the channels' existence.[9]

3.2 Asymptotic Costs for Small Discount Rates

The next proposition obtains closed-form formulas for the minimal costs of a unidirectional channel and a symmetric bidirectional channel. The general, asymmetric bidirectional channel is discussed separately.

Theorem 3.4 (Asymptotic Channel Costs)**.** *In the limit of r near zero:*

(i) The minimal cost of a unidirectional channel with rate λ is

$$L^{opt}(0,\lambda) = 2\left(\frac{B\lambda}{r}\right)^{1/2} - \frac{B}{2} + O(r^{1/2}) \tag{7}$$

and is achieved for the channel size $l_2 = \left(\frac{B\lambda}{r}\right)^{1/2} + O(r^{1/2})$.

(ii) The minimal cost of a symmetric bidirectional channel with equal rates λ is

$$L^{opt}(\lambda,\lambda) = 3\left(\frac{2B\lambda}{r}\right)^{1/3} - \frac{B}{6} + O(r^{1/3}) \tag{8}$$

and is achieved for channel sizes $l_1 = l_2 = \left(\frac{2B\lambda}{r}\right)^{1/3} + O(r^{1/3})$.

(iii) A nearly-symmetric bidirectional channel with $\lambda_2 - \lambda_1 = O(r^{1/3})$ has minimal cost

$$L^{opt}(\lambda_1,\lambda_2) = \left(3 + \frac{1}{2}\frac{\lambda_2-\lambda_1}{\lambda_1}\right)\left(\frac{2B\lambda_1}{r}\right)^{1/3} - \frac{B}{6} + O(r^{1/3}) \tag{9}$$

with the minimal channel sizes

$$l_1 = \left(\frac{2B\lambda_1}{r}\right)^{1/3} - \frac{\lambda_2-\lambda_1}{6\lambda_1}\left(\frac{2B\lambda_1}{r}\right)^{2/3} + O(r^{1/3}), \tag{10}$$

$$l_2 = \left(\frac{2B\lambda_1}{r}\right)^{1/3} + \frac{\lambda_2-\lambda_1}{6\lambda_1}\left(\frac{2B\lambda_1}{r}\right)^{2/3} + O(r^{1/3}). \tag{11}$$

[9]Note also that the conditions in Theorem 3.3 are necessary for optimally establishing a channel, but not sufficient. Nevertheless, necessity is all that is required to infer that r needs to be small relative to payment rates for all channels that should not be closed to reduce costs.

The main message of the above proposition is that both the minimal cost of a unidirectional channel and its required capital are of the order of $r^{-1/2}$. By contrast, for a symmetric bidirectional channel both the minimal cost and required capitals are of the order of $r^{-1/3}$. Note that, while in a unidirectional channel only the paying party commits collateral, a symmetric bidirectional channel requires both parties to commit collateral, but such collateral declines more slowly as r approaches zero.

For the general case of an asymmetric bidirectional channel with significantly different $\lambda_1 < \lambda_2$, the situation is more complex. The following result is obtained under the assumption that the balances l_1 and l_2 satisfy specific asymptotic properties in r, which are motivated by numerical calculations of the optimal $l_1(r), l_2(r)$ for smaller and smaller values of r, suggesting that the commitment of the average payer should be $l_2 = O(r^{-1/2})$, while the commitment of the average payee should be $l_1 = O(\log(r^{-1}))$.

Theorem 3.5 (Asymmetric Bidirectional Cost)**.** *If $l_2 = O(r^{-1/2})$ and $l_1 = O(\log(r^{-1}))$, then the minimal cost is*

$$2\left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2} + \frac{1+\log\left(1+\left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2}\log\frac{\lambda_2}{\lambda_1}\right)}{\log\frac{\lambda_2}{\lambda_1}} + O(1) \tag{12}$$

and the corresponding optimal channel sizes are

$$l_1 = \frac{\log\left(1+\left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2}\log\frac{\lambda_2}{\lambda_1}\right)}{\log\frac{\lambda_2}{\lambda_1}} + O(1), \tag{13}$$

$$l_2 = \left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2} + O(1). \tag{14}$$

The main message of this result is that, for r small enough, an asymmetric bidirectional channel is more akin to a unidirectional channel than to a symmetric bidirectional channel: both its minimal cost and the total required capital are of the order of $r^{-1/2}$, as in the unidirectional case, rather than of order $r^{-1/3}$, as in the symmetric bidirectional case. At the leading order, an asymmetric bidirectional channel is equivalent to a unidirectional channel with size $\lambda_2 - \lambda_1$, thereby considering only the overall net flow of transactions. Likewise, the amount of capital l_2 that the average payer (node 2) has to commit is the same as for a unidirectional channel of size $\lambda_2 - \lambda_1$.

Note that such approximate equivalence hinges on a rather delicate choice of the capital committed by the average payee (node 1), which is neither of order $r^{-1/2}$ or $r^{-1/3}$, but of the much lower order $\log r^{-1}$. Thus, while the average payer benefits from the bidirectional channel only through the netting effect, the average payee is the main beneficiary of the arrangement, by committing an amount that is logarithmic in the capital committed by the average payer. For example, if one node pays the other monthly ($\lambda_1 = 12$), while the other pays weekly ($\lambda_2 = 52$), with a reset cost of $B = 0.1$ and a discount rate $r = 1\%$, the average payer commits $l_2 \approx 20$ to the channel, while the average payee commits only $l_1 \approx 3$. Reducing the interest rate to $r = 0.1\%$, the difference becomes even starker, with l_2 exceeding 63 while l_1 barely increasing to 3.8.

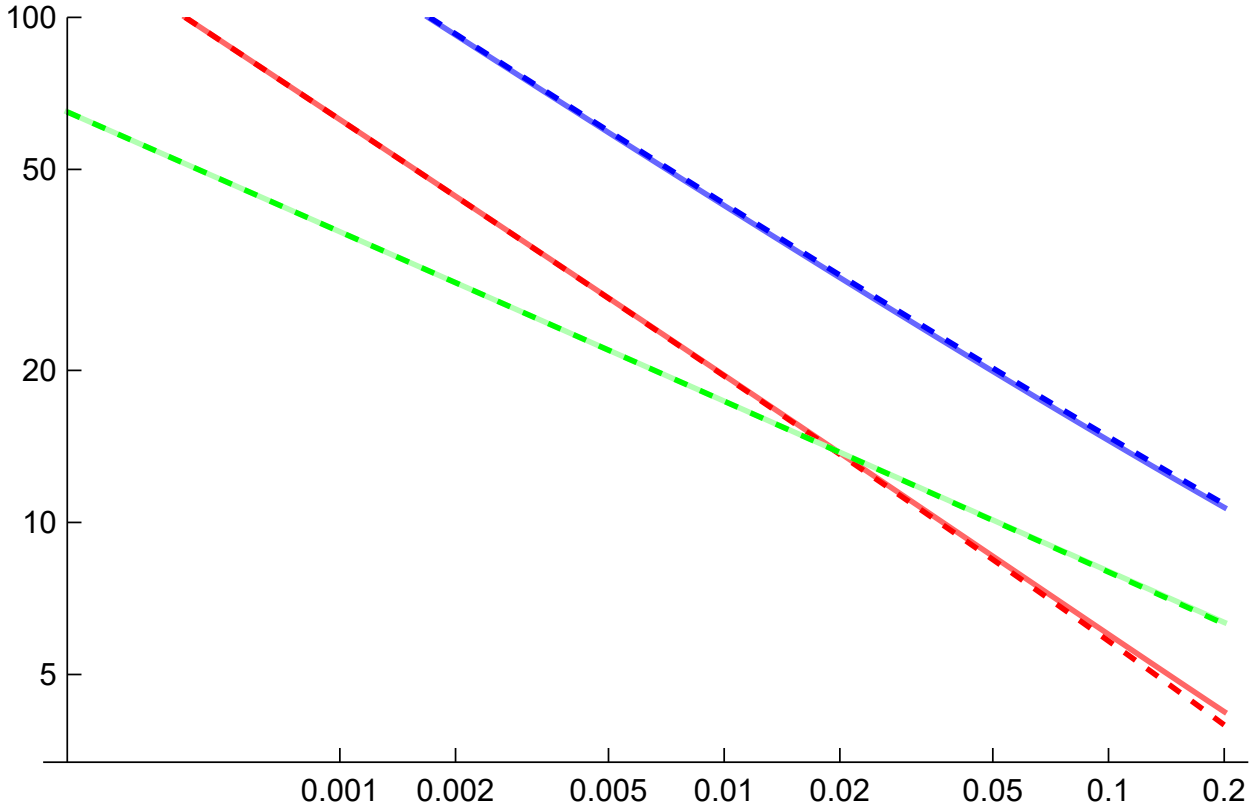

Figure 4: Exact (solid) and asymptotic (dashed) cost (vertical) of unidirectional (red, (2) vs. (7)), bidirectional symmetric (green, minimum of (3) vs. (8)), and bidirectional asymmetric (blue, minimum of (3) vs. (13)) channels against the interest rate (horizontal), for $B = 1, \lambda_1 = 1$, and (for the blue plot) $\lambda_2 = 5$. Both axes are in logarithmic scale.

The asymptotic formulas in Proposition 3.4 (iii) and Proposition 3.5 offer different approximations of the optimal l_1 and l_2 for given $\lambda_1, \lambda_2, B, r$. Both approximations become increasingly accurate as $r \downarrow 0$, with (10) becoming superior to (13) when $(\lambda_2 - \lambda_1)/r^{1/3}$ is relatively small, and vice versa when $(\lambda_2 - \lambda_1)/r^{1/3}$ is relatively large, which corresponds to λ_1, λ_2 fixed as $r \downarrow 0$.

The asymptotic approximations are very accurate for the typical range of the interest rate: as shown by Figure 4, exact formulas are virtually indistinguishable from their asymptotic approximations for rates below 20%, with minor deviations arising only in the unidirectional approximation, and for rates close to 20%.

The next result calculates the probability that either node exhausts the balance before the other. This problem is trivial for unidirectional and symmetric channels, but the general case is more delicate: in the nearly-symmetric regime the probabilities of exhaustion are different and nontrivial, while in the asymmetric regime the probability that the average payee exhausts the balance before the average payer is negligible.

Proposition 3.6 (Probability of balance exhaustion)**.** *Let* $q = \lambda_2/\lambda_1 \geq 1$.

(i) In a channel with sizes l_1, l_2 *and rates* λ_1, λ_2*, the probability that node 2 exhausts the*

channel before node 1 is (a) $\frac{1-q^{-l_1}}{1-q^{-l_1-l_2}}$ *if* $\lambda_2 > \lambda_1 > 0$*, (b)* $1/2$ *if* $\lambda_2 = \lambda_1 > 0$*, and (c) 1 if* $\lambda_2 > \lambda_1 = 0$*.*

(ii) Thus, under the optimal choice of l_1 *and* l_2*, such probability is identically one for a unidirectional channel (Theorem 3.4(i)) and equals to 1/2 for a symmetric channel (Theorem 3.4(ii)).*

(iii) As $r \to 0$ *and under the optimal choice of* l_1 *and* l_2*, such probability converges to one for a bidirectional channel as in Proposition 3.5. For a nearly-symmetric channel (Theorem 3.4(iii)), it converges to* $\frac{z - z^{\frac{1}{24}\log z + \frac{1}{2}}}{z-1}$*, where* $z = \exp\left(\frac{2^{4/3}B^{1/3}(\lambda_2-\lambda_1)}{\lambda_1^{2/3}r^{1/3}}\right)$*.*

Channel deployment reduces the frequency of on-chain transactions, thereby reducing the congestion of the queue to transact on-chain. The next theorem provides an asymptotic formula for the *on-chain rate*, that is, the number of average channel transactions for which an on-chain transaction is required. Such a ratio obviously depends on the channel sizes l_1, l_2. It also depends on the transaction rates λ_1, λ_2, but only through their ratio.

Proposition 3.7 (Congestion Reduction)**.** *In a bidirectional channel with sizes* l_1, l_2 *and rates* λ_1, λ_2*, the long-term ratio between the number of channel transactions and the number of on-chain transactions equals*

$$\frac{(q+1)\left(\left(l_2\left(q^{l_1}-1\right)-l_1\right)q^{l_2}+l_1\right)}{(q-1)\left(q^{l_1+l_2}-1\right)}$$

where $q = \lambda_2/\lambda_1$*. In particular:*

(i) In a unidirectional channel ($q \downarrow 0$) *the ratio simplifies to* l_1*.*

(ii) In a symmetric bidirectional channel ($q \to 1$) *the ratio simplifies to* $l_1 l_2$*.*

The above proposition shows that, while a unidirectional channel reduces on-chain traffic in proportion to its size, for a symmetric bidirectional channel the reduction is proportional to the product of the two sizes, hence much more significant for a total committed amount. In particular, when the balances committed are set optimally, (i) the reduction is higher for those nodes from which incoming payments exceed outgoing payments, and (ii) the reduction is highest for symmetric channels.

As the numerical results in Figure 1 show, if both payment rates are large enough, the single bidirectional channel becomes optimal and absorbs both payment flows. In particular, it is never optimal to use two separate unidirectional channels. This conclusion, however, does not mean that two unidirectional channels are always more costly than the corresponding bidirectional channel, and a close inspection reveals that there are cases in which they may be cheaper. But the point is that, in such cases, it is even cheaper for one payment flow to take place on chain, thereby excluding two unidirectional channels from the optimal configurations.

4 Discussion

This paper formulates and analyzes a parsimonious model of bitcoin transactions that can be shifted to a bidirectional LN channel. The present analysis covers their costs, benefits and circumstances in which they are useful. A follow-up paper (Guasoni et al., 2021) explores the implications to the topology of the LN.

The model analyzed here takes the on-chain channel reopening transaction fee B as fixed and exogenous to the model. Earlier work (Huberman et al., 2021, 2019) argues that Bitcoin payment system users pay these fees when the system is congested. When users' delay costs vary, so will the fees they offer. The fees do not affect the protocol-determined throughput of the Bitcoin payment system but they induce the miners to assign processing priority to the transactions associated with higher fees.

The present paper shows that each channel reduces the number of on-chain transactions. Thereby the LN as a whole reduces on-chain congestion and the fee C. Future work will model the interaction between the level of the transaction fee, the level of congestion and the throughput improvement due to the availability of a layer-two solution such as the LN.

The more transactions are shifted from the blockchain to the LN, the more beneficial the LN. In a unidirectional channel to which l units are committed, there are l in-channel transactions for each on-chain transaction. In a symmetric bidirectional channel to which l units are committed ($l/2$ on each side), there are $l^2/4$ in-channel transactions for each on-chain transaction. In general, the closer a bidirectional channel is to symmetric, the more economically beneficial it is.

Successful layer-two solutions are practical, convenient, often render low transaction size economically feasible, support higher transaction throughput than the first layer, and have the potential to gradually diminish the role of the underlying first layer.

Famously, the Bitcoin payment network has low throughput and high latency. The LN is a payment solution built on top of the Bitcoin payment network designed to address these weaknesses. If Bitcoin becomes popular, it is likely that so will be the LN or a future version of it. Moreover, the LN and its relation to Bitcoin serve as prototypes and proofs of concepts for future payment systems and therefore are study-worthy.

5 Proofs

Proof. Proof of Lemma 3.1 Denote by τ_1 the arrival time of the next transaction, which is an exponential random variable with rate λ (likewise, the future arrival times are denoted by τ_n for $n \geq 1$). As the interarrival times of future transactions are independent of previous arrival times, the expected cost $\mu = E[\sum_{n=1}^{\infty} e^{-r\tau_n} C]$ of the entire stream satisfies, in view of the Markov property of the Poisson process,

$$\mu = E\left[e^{-r\tau_1}\left(C + E\left[\sum_{n=1}^{\infty} e^{-r(\tau_n-\tau_1)} C \Big| \tau_1\right]\right)\right] = E\left[e^{-r\tau_1}\left(C+\mu\right)\right]$$

and hence $\mu = \frac{\lambda}{\lambda+r}(B/2+\mu)$ which in turn implies that $\mu = \frac{\lambda B}{2r}$. □

Proof. Proof of Theorem 3.2 Let X_t denote the net cumulative balance at time t of node 1, from both stream of transactions, i.e.,

$$X_t = N_t^{\lambda_2} - N_t^{\lambda_1}.$$

Thus, an increase in X represents money flowing to node 1, a decrease money flowing to node 2. The cash balance starts at zero, and varies over time according the dynamics of the Poisson processes. Node 1 commits l_1 to the channel and node 2 commits l_2.

Denote by $J(n)$ the expected total future cost when the balance at time t is equal to n, with $-l_1 < n < l_2$, and by τ the time that elapses from t until the next transaction. Because $\tau = \tau^{\lambda_2} \wedge \tau^{\lambda_1}$ is the minimum between two independent exponential random variables τ^{λ_2} and τ^{λ_1} with rates λ_2 and λ_1 respectively, it is also an exponential variable with rate $\lambda_2 + \lambda_1$. When τ arrives, $\tau = \tau^{\lambda_2}$ with probability $\lambda_2/(\lambda_2 + \lambda_1)$, in which case $X_{t+\tau} = X_t + 1$. Otherwise, $\tau = \tau^{\lambda_1}$ with probability $\lambda_1/(\lambda_2 + \lambda_1)$, and hence $X_{t+\tau} = X_t - 1$. Thus, the expected cost $J(n)$ satisfies the equation

$$J(n) = E\left[\int_0^\tau e^{-rs} r(l_1 + l_2) ds\right] + E[e^{-r\tau} J(X_{t+\tau}) | X_t = n],$$

where the first term represents the opportunity cost in the time interval $[t, t+\tau]$ and the second term the residual expected cost from τ onwards. Note that $E\left[e^{-r\tau}\right] = (\lambda_2 + \lambda_1)/(\lambda_2 + \lambda_1 + r)$ because τ is exponentially distributed with rate $\lambda_2 + \lambda_1$. Thus, the first term equals

$$E\left[\int_0^\tau e^{-rs} r(l_1 + l_2) ds\right] = (l_1 + l_2) E\left[(1 - e^{-r\tau})\right] =$$
$$= (l_1 + l_2)(1 - E\left[e^{-r\tau}\right]) = (l_1 + l_2)\left(1 - \frac{\lambda_2 + \lambda_1}{r + \lambda_2 + \lambda_1}\right) = (l_1 + l_2)\frac{r}{r + \lambda_2 + \lambda_1}$$

while the last term equals

$$\begin{aligned}
E[e^{-r\tau} J(X_{t+\tau}) | X_t = n] &= E[e^{-r\tau} J(X_{t+\tau}) | X_t = n, \tau = \tau^{\lambda_2}] P(\tau = \tau^{\lambda_2}) \\
&+ E[e^{-r\tau} J(X_{t+\tau}) | X_t = n, \tau = \tau^{\lambda_1}] P(\tau = \tau^{\lambda_1}) \\
&= J(n+1) E[e^{-r\tau}] \frac{\lambda_2}{\lambda_2 + \lambda_1} + J(n-1) E[e^{-r\tau}] \frac{\lambda_1}{\lambda_2 + \lambda_1} \\
&= J(n+1) \frac{\lambda_2 + \lambda_1}{r + \lambda_2 + \lambda_1} \frac{\lambda_2}{\lambda_2 + \lambda_1} + J(n-1) \frac{\lambda_2 + \lambda_1}{r + \lambda_2 + \lambda_1} \frac{\lambda_1}{\lambda_2 + \lambda_1} \\
&= J(n+1) \frac{\lambda_2}{r + \lambda_2 + \lambda_1} + J(n-1) \frac{\lambda_1}{r + \lambda_2 + \lambda_1}
\end{aligned}$$

In short, the expected cost function satisfies the difference equation

$$J(n) = (l_1 + l_2)\frac{r}{r + \lambda_2 + \lambda_1} + J(n+1)\frac{\lambda_2}{r + \lambda_2 + \lambda_1} + J(n-1)\frac{\lambda_1}{r + \lambda_2 + \lambda_1}.$$

Bidirectional channel. In a bidirectional channel, such difference equation is combined with the boundary conditions

$$J(-l_1) = J(0) + B \qquad J(+l_2) = J(0) + B,$$

which require that, once a liquidation point is reached, the residual expected cost equals the expected reset cost B plus future costs starting from the reset state 0. The general solution to (5) is

$$J(n) = l_1 + l_2 + k_1\alpha_-^n + k_2\alpha_+^n$$

where $\alpha_\pm$ are as in (4). Substituting the general form of $J(n)$ into the boundary conditions, one obtains two linear equations for k_1 and k_2, which yield the cost function:

$$J(n) = l_1 + l_2 - \frac{B\left(\alpha_-^{l_1+l_2} - 1\right)\alpha_+^{l+n}}{\alpha_-^{l_1} - \alpha_+^{l_1} + \left(\alpha_+^{l_1} - 1\right)\alpha_-^{l_1+l_2} - \left(\alpha_-^{l_1} - 1\right)\alpha_+^{l_1+l_2}} - \frac{B\left(\alpha_+^{l_1+l_2} - 1\right)\alpha_-^{l+n}}{\alpha_+^{l_1} + \alpha_-^{l_1}\left(-\left(\alpha_+^{l_1} - 1\right)\alpha_-^{l_2} - 1\right) + \left(\alpha_-^{l_1} - 1\right)\alpha_+^{l_1+l_2}}$$

from which in turn the initial cost $J(0)$ in (3) follows.

Unidirectional channel. The cost of a unidirectional channel follows from a similar argument: as only one node pays the other (suppose that only 1 is paid, whence $\lambda_1 = 0$), it follows that the other party needs not to commit capital ($l_1 = 0$), whence

$$J(n) = l_2\frac{r}{r+\lambda_2} + J(n+1)\frac{\lambda_2}{r+\lambda_2}$$

which has the general solution

$$J(n) = l_2 + k\left(\frac{r+\lambda_2}{\lambda_2}\right)^n$$

where the constant k is determined by the boundary condition

$$J(l_2) = J(0) + B.$$

Thus, the cost function is

$$J(n) = l_2 + B\left(\left(\frac{r+\lambda_2}{\lambda_2}\right)^{l_2} - 1\right)^{-1}\left(\frac{r+\lambda_2}{\lambda_2}\right)^n$$

whence the initial cost in (1). The minimal formula follows by differentiating the above formula with respect to l_2, solving for the value of l_2 for which the derivative is null, and replacing the resulting value in the formula itself. □

Proof. Proof of Theorem 3.3 (i) By Theorem 3.2 (i), the cost of a unidirectional channel of size m is

$$m + B\left(\left(1 + \frac{r}{\lambda}\right)^m - 1\right)^{-1}.$$

To ascertain whether it is worth to establish such a channel, one needs to compare such cost with the alternative of establishing no channel at all, which is $\lambda C/r$ by Lemma 3.1. Thus, a one-directional channel is suboptimal if and only if

$$\inf_{m>0}\left(m + B\left(\left(1 + \frac{r}{\lambda}\right)^m - 1\right)^{-1} - \frac{\lambda C}{r}\right) > 0.$$

To find such an infimum, denote by $\lambda/r = (e^\kappa - 1)^{-1}$, which allows to rewrite the function to minimize as

$$F(m) = \frac{C}{1-e^\kappa} + \frac{B}{e^{\kappa m} - 1} + m.$$

It is immediate to see that such a function (i) is convex, and (ii) diverges to $+\infty$ as m approaches 0 or ∞. Thus, the function admits a unique minimum for $\hat{m} \in (0, \infty)$, and such minimum satisfies the first-order condition $F'(\hat{m}) = 0$, i.e.,

$$\frac{B\kappa}{2 - 2\cosh(\kappa m)} + 1 = 0$$

which yields

$$\hat{m} = \frac{\cosh^{-1}\left(\frac{B\kappa}{2} + 1\right)}{\kappa} \text{ and } F(\hat{m}) = \frac{C}{1-e^\kappa} + \frac{2B}{B\kappa + \sqrt{B\kappa(B\kappa+4)}} + \frac{\cosh^{-1}\left(\frac{B\kappa}{2} + 1\right)}{\kappa}.$$

Hence, the channel is suboptimal if and only if $F(\hat{m}) > 0$. Thus, to obtain a sufficient condition for this property, it is enough to find a lower bound for $F(\hat{m})$ and require that it is positive. For this purpose, note first that the elementary estimate $e^x > 1 + x$ for $x > 0$ implies that $C/(1-e^\kappa) > -C/\kappa$ for all $\kappa > 0$. Note also that

$$\frac{2B}{B\kappa + \sqrt{B\kappa(B\kappa+4)}} \geq \frac{\alpha}{\kappa} \quad \text{for} \quad \kappa \geq \frac{\alpha^2}{B(1-\alpha)},\ 0 < \alpha < 1, \tag{15}$$

whence the lower bound

$$F(\hat{m}) > -\frac{C}{\kappa} + \frac{\alpha}{\kappa} + \frac{\cosh^{-1}\left(\frac{B\kappa}{2} + 1\right)}{\kappa} \quad \text{for} \quad \kappa \geq \frac{\alpha^2}{B(1-\alpha)}. \tag{16}$$

The lower bound is in turn positive for $\kappa \geq \frac{2}{B}(\cosh(C-\alpha) - 1)$, whence

$$F(\hat{m}) > 0 \quad \text{for} \quad \kappa \geq \frac{1}{B}\max\left(\frac{\alpha^2}{1-\alpha}, 2(\cosh(C-\alpha) - 1)\right) \tag{17}$$

Recall now that $\cosh(x) = (e^x + e^{-x})/2 \leq (e^{|x|} + 1)/2$, whence

$$2(\cosh(C-\alpha) - 1) \leq e^{|C-\alpha|} - 1 \tag{18}$$

Thus, if $\kappa \geq (e^C - 1)/B$, then $\kappa \geq \frac{2}{B}(\cosh(C-\alpha) - 1)$ and, choosing $\alpha \leq \frac{C}{C+1}$,

$$\kappa \geq \frac{e^C - 1}{B} \geq \frac{C}{B} \geq \frac{\alpha^2}{B(1-\alpha)} \tag{19}$$

whence $F(\hat{m}) \geq 0$ by (16). Because $\kappa = \log(1 + r/\lambda)$, the condition $\kappa \geq (e^C - 1)/B$ is equivalent to $r/\lambda \geq e^{(e^C-1)/B} - 1$. As this condition is sufficient for on-chain transactions to be cheaper than a unidirectional channel, it follows that a necessary condition for this channel's optimal existence is that $\lambda > r(e^{(e^C-1)/B} - 1)^{-1}$, which is equivalent to the claim. □

Proof. Proof of Theorem 3.3 (ii) Setting $\lambda_2 = \lambda_1 = \lambda$ in Theorem 3.2 (ii), it follows that the cost of a symmetric bidirectional channel is

$$l_1 + l_2 + \frac{B\left(\alpha^{l_1} + \alpha^{l_2}\right)}{\left(\alpha^{l_1} - 1\right)\left(\alpha^{l_2} - 1\right)}$$

where $\alpha := \alpha_+ = \frac{r+2\lambda+\sqrt{r(4\lambda+r)}}{2\lambda} > 1$. Subtracting from such cost the cost of doing on-chain transactions instead, which is $2C\lambda/r$, and rewriting λ/r in terms of α, the difference is

$$l_1 + l_2 + \frac{B\left(\alpha^{l_1} + \alpha^{l_2}\right)}{\left(\alpha^{l_1} - 1\right)\left(\alpha^{l_2} - 1\right)} - \frac{2C\alpha}{(\alpha-1)^2}$$

It is easy to check that this function is strictly convex, as the trace and the determinant of its Hessian are both positive. Thus, its minimizer must be unique. Because the function is symmetric in l_1 and l_2, the minimum must be achieved for $l_1 = l_2$, otherwise it would not be unique (if (l_1, l_2) were a minimizer, then (l_2, l_1) would be another minimizer). Thus, it suffices to minimize the above function for $l_1 = l_2 = m$, i.e.,

$$2m + \frac{2B\alpha^m}{\left(\alpha^m - 1\right)^2} - \frac{2C\alpha}{(\alpha-1)^2} \tag{20}$$

Setting $\alpha = e^\kappa$, the function reduces to

$$2m + \frac{B}{2}\operatorname{csch}^2\left(\frac{\kappa m}{2}\right) - \frac{C}{2}\operatorname{csch}^2\left(\frac{\kappa}{2}\right)$$

Minimizing this function is cumbersome, in that the minimum does not have a simple explicit solution. To find a tractable lower bound, recall the inequalities

$$\frac{1}{x^2} - \frac{1}{3} < \operatorname{csch}^2 x < \frac{1}{x^2}, \qquad x > 0, \tag{21}$$

whereby (20) is bounded from below by

$$2m + B\left(\frac{2}{\kappa^2 m^2} - \frac{1}{6}\right) - \frac{2C}{\kappa^2}. \tag{22}$$

For $\hat{m} = (2B/\kappa^2)^{1/3}$, this function reaches its minimum, which is

$$3\left(\frac{2B}{\kappa^2}\right)^{1/3} - \frac{B}{6} - \frac{2C}{\kappa^2}. \tag{23}$$

Thus, this quantity is positive if and only if $C \leq 3\left(\frac{B\kappa^4}{4}\right)^{1/3} - \frac{B}{12}\kappa^2$, which is thus a sufficient condition for the stream of transactions to be cheaper on-chain than over the channel. The reverse inequality is therefore a necessary condition for the channel to be cheaper. □

Proof. Proof of Theorem 3.4 (i) For small values of r, the cost in (1) simplifies to

$$J(0) = \frac{B\lambda}{rl_2} + \left(l_2 - \frac{B(l_2-1)}{2l_2}\right) + O(r)$$

which is maximized by

$$l_2 = \left(\frac{B\lambda}{r}\right)^{1/2} + O(r^{1/2}).$$

Plugging this formula into (5) in turn yields the minimal cost (7).

(ii) For a symmetric channel and a small interest rate ($r \downarrow 0$), the values of $\alpha_\pm$ in (4) simplify to

$$\alpha_\pm = 1 + \frac{r \pm \sqrt{r^2+4r\lambda}}{2\lambda} \approx 1 \pm \sqrt{\frac{r}{\lambda}} + \frac{r}{2\lambda} + O(r^{3/2})$$

and the value function becomes in turn, at order zero,

$$J(0) = \frac{2B\lambda}{l_1 l_2 r} + \frac{B\left(l_1^2 - 3l_1l_2 + l_2^2 + 1\right)}{6l_1l_2} + l_1 + l_2 + O(r^{1/2}).$$

The advantage of this expression is that its minimizers l_1 and l_2 can be found explicitly. Indeed, the first order conditions for l_1 and l_2 are respectively

$$1 - \frac{B\left(12\lambda + r\left(-l_1^2 + l_2^2 + 1\right)\right)}{6l_1^2 l_2 r} = 0 \qquad 1 - \frac{B\left(12\lambda + r\left(l_1^2 - l_2^2 + 1\right)\right)}{6l_1 l_2^2 r} = 0$$

And the (real) solution to this system is, at the leading order,

$$l_1 = l_2 = \left(\frac{2B\lambda}{r}\right)^{1/3} + O(r^{1/3})$$

Substituting this expression into the objective function yields the minimal cost in (8).

(iii) Up to a subsequence, assume that $\lambda_2(r) \sim \lambda_1 + kr^{1/3}$ for some constant $k > 0$. Then the asymptotic expansion of the value function, at the first order in k and at the zero-order in r, is

$$\begin{aligned} J(0) =& \frac{2B\lambda_1}{l_1 l_2 r} + \frac{2Bk(l_1 - l_2 + 3)}{6l_1 l_2 r^{2/3}} + \\ &+ \frac{l_1^2(B + 6l_2) + 3l_2 l_1(2l_2 - B) + Bl_2^2 + B}{6l_1 l_2} + O(k^2 r^{1/3}) \end{aligned} \tag{24}$$

Maximizing this objective with respect to l_1 and l_2 yields the first-order conditions

$$\begin{aligned} \frac{Bk(l_2 - 3)r^{1/3}}{3l_1^2 l_2} - \frac{2B\lambda_1}{l_1^2 l_2} + r\left(-\frac{Bl_2}{6l_1^2} - \frac{B}{6l_1^2 l_2} + \frac{B}{6l_2} + 1\right) = 0 \\ -\frac{Bk(l_1 + 3)r^{1/3}}{3l_1 l_2^2} - \frac{2B\lambda_1}{l_1 l_2^2} + r\left(\frac{B}{6l_1} - \frac{Bl_1}{6l_2^2} - \frac{B}{6l_1 l_2^2} + 1\right) = 0 \end{aligned}$$

whose solutions are, at the same order:

$$l_1 = \left(\frac{2B\lambda_1}{r}\right)^{1/3} - \frac{B^{2/3}k}{3(2\lambda_1 r)^{1/3}}, \qquad l_2 = \left(\frac{2B\lambda_1}{r}\right)^{1/3} + \frac{B^{2/3}k}{3(2\lambda_1 r)^{1/3}}$$

and replacing $k = (\lambda_2 - \lambda_1)/r^{1/3}$ in the above expressions and in (24), the formulas in (9) follow. □

Proof. Proof of Theorem 3.5 If $l_2(r) = O(r^{-1/2})$ and $l_1(r) = O(\log r^{-1})$, then there exists $\zeta_1, \zeta_2 \in \mathbb{R}$ and a sequence $(r_k)_{k\geq 1}$, decreasing to zero, such that $\lim_{k\to\infty} l_1(r_k)/\log r_k = \zeta_1$ and $\lim_{k\to\infty} l_2(r_k)/r_k^{1/2} = \zeta_2$. Along such a sequence, from (3) and (4) it follows that:

$$\lim_{k\to\infty} (L^{l_1(r_k),l_2(r_k)}(\lambda_1(r_k),\lambda_2(r_k)) - l_1(r_k) - l_2(r_k))r_k^{1/2} = B\frac{\lambda_2-\lambda_1}{\zeta_2}.$$

Therefore, for such subsequence the cost equals:

$$\left(B\frac{\lambda_2-\lambda_1}{\zeta_2} + \zeta_2\right) r_k^{-1/2}.$$

Thus, the only value of ζ_2 for which $l_2(r)$ can be optimal must be the minimizer of this expression. Minimizing it with respect to ζ_2 yields the minimizer $\hat{\zeta}_2 = (B(\lambda_2-\lambda_1))^{1/2}$ and the minimum $2(B(\lambda_2-\lambda_1)/r)^{1/2}$. In particular, such $\hat{\zeta}_2$ is optimal for any subsequence, and therefore $\lim_{r\to 0} l_2(r)/r^{1/2} = \hat{\zeta}_2$. Likewise, to calculate the second-order term, it suffices to calculate the expansion of

$$L^{l_1(r),l_2(r)}(\lambda_1(r),\lambda_2(r)) - l_1(r) - l_2(r)$$

for $l_2(r) = \left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2} + O(1)$ and, on the subsequence considered, such an expansion equals

$$\left(1 - (\lambda_2/\lambda_1)^{-\zeta_1}\right)^{-1} \left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2}$$

which entails that the second-order term of the total cost is

$$\zeta_1 \log r + \left(1 - (\lambda_2/\lambda_1)^{-\zeta_1}\right)^{-1} \left(\frac{B(\lambda_2-\lambda_1)}{r}\right)^{1/2} + O(1) \tag{25}$$

Minimizing this expression over ζ_1, one obtains the minimizer

$$\hat{\zeta}_1 = \frac{\log\left(\log\left(\frac{\lambda_2}{\lambda_1}\right)\sqrt{\frac{B(\lambda_2-\lambda_1)}{r}} + 1\right)}{\log\left(\frac{\lambda_2}{\lambda_1}\right)} + O(1)$$

which must hold for any subsequence, and substituting it into (25) yields (12). □

Proof. Proof of Proposition 3.6 If $\lambda_2 > \lambda_1 = 0$, then X_t is increasing, therefore the claim is trivial. If $\lambda_2 = \lambda_1 > 0$, then the balance X_t is a martingale, and its probability of reaching l_2 before reaching $-l_1$ is $l_1/(l_1 + l_2)$, which is $1/2$ for $l_1 = l_2$.

If $\lambda_2 > \lambda_1 > 0$, denote by $p(n)$ the probability that node 2 liquidates before node 1, if the current balance is n. Thus, by definition, $p(l_2) = 1$ and $p(-l_1) = 0$. Because the balance moves from n to $n+1$ with probability $\lambda_2/(\lambda_1 + \lambda_2)$ and to $n-1$ with probability $\lambda_1/(\lambda_1 + \lambda_2)$, $p(n)$ satisfies

$$p(n) = \frac{\lambda_1}{\lambda_1 + \lambda_2} p(n-1) + \frac{\lambda_2}{\lambda_1 + \lambda_2} p(n+1)$$

and the general solution of this difference equation is $p(n) = a + bq^{-n}$, where $q = \lambda_2/\lambda_1$. The constants a and b are identified by the conditions $p(l_2) = 1$, $p(-l_1) = 0$, whence

$$p(n) = \frac{1 - q^{-n-l_1}}{1 - q^{-l_1-l_2}}$$

which reduces to the claim for $n = 0$.

As $r \to 0$, note that in the setting of Proposition 3.5 both λ_1, λ_2 remain fixed, while $l_1(r), l_2(r)$ diverge to infinity, therefore $p(0)$ converges to 1. Vice versa, in the nearly symmetric case (Theorem 3.4(iii)), $\lambda_1(r), \lambda_2(r)$ converge to 1, while $l_1(r), l_2(r)$ diverge to infinity, and the corresponding probability follows by taking the limit of the resulting expression as $r \to 0$. □

Proof. Proof of Proposition 3.7 Let $m(n)$ be the fraction of transactions that the balance X spends in state $-l_1 \leq n \leq l_2$. Consider first n positive: for $0 < n < l_2$, $m(n)$ satisfies

$$m(n) = \frac{\lambda_2}{\lambda_1 + \lambda_2} m(n-1) + \frac{\lambda_1}{\lambda_1 + \lambda_2} m(n+1)$$

because state n can only be reached from either $n-1$ through an up-move, which has probability $\lambda_2/(\lambda_1 + \lambda_2)$, or from $n+1$ through a down-move, which has probability $\lambda_1/(\lambda_1 + \lambda_2)$. By construction, once the state l_2 is reached, the state is immediately reset to 0, hence $m(l_2) = 0$ and

$$m(n) = m(0) \frac{q^{l_2} - q^n}{q^{l_2} - 1} \qquad 0 \leq n \leq l_2. \tag{26}$$

Likewise, note that (5) is also valid for $-l_1 < n < 0$ and that $m(-l_1) = 0$ by construction, whence

$$m(n) = m(0) \frac{q^{l_1+n} - 1}{q^{l_1} - 1} \qquad -l_1 \leq n \leq 0. \tag{27}$$

Then, the condition $\sum_{n=-l_1}^{l_2} m(n) = 1$ yields a linear equation that identifies

$$m(0) = \frac{\left(q^{l_1} - 1\right)\left(q^{l_2} - 1\right)}{\left(l_2\left(q^{l_1} - 1\right) - l_1\right) q^{l_2} + l_1}. \tag{28}$$

Finally, note that the channel transactions that lead to an on-chain transaction are the up-moves from $l_2 - 1$ and the down-moves from $-l_1 + 1$. Thus, the fraction of such transactions is

$$m(l_2 - 1) \frac{\lambda_2}{\lambda_1 + \lambda_2} + m(-l_1 + 1) \frac{\lambda_1}{\lambda_1 + \lambda_2}$$

and its value is obtained from (26), (27), and (28). The reciprocal of such value is precisely the long-term average of number of channel transactions per on-chain transaction in (3.7). □

References

Allais, M. (1947), *Economie & interet: presentation nouvelle des problemes fondamentaux relatifs au role economique du taux de l'interet et de leurs solutions*, Librairie des publications officielles.

Baumol, W. J. (1952), 'The transactions demand for cash: An inventory theoretic approach', *The Quarterly Journal of Economics* pp. 545–556.

Brânzei, S., Segal-Halevi, E. and Zohar, A. (2022), How to charge lightning: The economics of bitcoin transaction channels, *in* '2022 58th Annual Allerton Conference on Communication, Control, and Computing (Allerton)', IEEE, pp. 1–8.

Frost, J., Shin, H. S. and Wierts, P. (2020), An early stablecoin? the bank of amsterdam and the governance of money, Technical report.

Guasoni, P., Huberman, G. and Shikhelman, C. (2021), Lightning network economics: Topology.

Huberman, G., Leshno, J. D. and Moallemi, C. (2019), An economist's perspective on the bitcoin payment system, *in* 'AEA Papers and Proceedings', Vol. 109, pp. 93–96.

Huberman, G., Leshno, J. D. and Moallemi, C. (2021), 'Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System', *The Review of Economic Studies* . rdab014.
URL: *https://doi.org/10.1093/restud/rdab014*

Ibaba, A., Abdulmalik, A., Alhassan, A.-B., Sunusi, I. and Thomas, S. (2021), Rem9ja: Bitcoin lightning network remittance solution, *in* '2021 1st International Conference on Multidisciplinary Engineering and Applied Science (ICMEAS)', IEEE, pp. 1–4.

Miller, M. H. and Orr, D. (1966), 'A model of the demand for money by firms', *The Quarterly journal of economics* **80**(3), 413–435.

Nakamoto, S. (2008), Bitcoin: A peer-to-peer electronic cash system, Technical report, Manubot.

Narayanan, A., Bonneau, J., Felten, E., Miller, A. and Goldfeder, S. (2016), *Bitcoin and cryptocurrency technologies: a comprehensive introduction*, Princeton University Press.

Poon, J. and Dryja, T. (2015), 'The bitcoin lightning network: Scalable off-chain instant payments'.

Tetek, J. (2021), 'Strike: Disrupting the remittance payments via Bitcoin's Lightning Network, url = https://vacuumlabs.com/strike-disrupting-remittance-payments/'.

Tobin, J. (1956), 'The interest-elasticity of transactions demand for cash', *The review of Economics and Statistics* pp. 241–247.

von Luckner, C. G., Reinhart, C. M. and Rogoff, K. S. (2021), Decrypting new age international capital flows, Technical report, National Bureau of Economic Research.

Wirdum, A. V. (2016), 'Understanding the lightning network'.
URL: *https://bitcoinmagazine.com/articles/understanding-the-lightning-network-part-building-a-bidirectional-payment-channel-1464710791*