Will Lightning Work

By Andrew Bakst

What's at stake

Bitcoin is a 100 billion dollar network banking on a relatively under-capitalized company to achieve scalability, and the stakes couldn't be higher.

If Bitcoin can't process a significant amount of transactions, does Bitcoin Cash win? Is there a new, even more contentious hard fork? Ripple, Dash, ZCash, Monero, Litecoin, and Bitcoin Cash investors are heavily incentivized to root for Lightning Network's failure, as are investors in "third generation" currencies such as Coda, Beam, and Grin. Even proponents of smart contract platforms are incentivized to see Lightning fail due to their competition with Bitcoin on becoming a global store of value. Becoming a better version of gold (in the near term) or an alternative global reserve currency (in the long term) offers the greatest total addressable market in the crypto space.

This analysis is an exploration of whether the Bitcoin blockchain will be able to achieve successful layer two scalability through the Lightning Network. Liquid by Blockstream and other side chain scaling solutions require immense trust in centralized parties, diminishing Bitcoin's greatest value proposition. Only if Bitcoin remains decentralized can black swan events like the Gold Reserve Act be avoided or at least fought against valiantly. Enter Lightning.

There are already plenty of resources explaining the Lightning Network:

For a short summary on Lightning;
For a more in depth technical explanation of how the Lightning Network works;

For Lightning to succeed at scale, hubs will need to exist. Without hubs, Lightning's usability is diminished significantly--routing across a plethora of nodes, while a solvable problem, requires increased costs to users (each node will take a fee) and carries a higher risk that a payment does not reach its destination address (a greater number of involved nodes means a greater likelihood of encountering an adversary).

While some argue that hubs make Bitcoin centralized, hubs do maintain Bitcoin's promise of being a bearer instrument: hubs cannot confiscate or lend users' funds unless a user explicitly grants the hub permission to do so. In side chain scaling solutions, users have no choice but to grant the side chain's validators permission to confiscate their Bitcoin. Hubs are a compromise between the benefits of decentralization (no additional parties have access to funds without opt-in user consent) and the benefits of centralization (transactions are fast and have low fees).

A Lightning hub will be forced to lock a significant amount of capital in the Lightning network. Deployed capital has an opportunity cost. The 'risk-free' (LIBOR) rate is currently 3%, and the borrow rate of Bitcoin is currently ~12%. 12% is hard to beat in any market right now, especially given that it can be obtained without any operational costs and with very low risk. Lenders currently require significant collateral in dollars from their borrowers, sometimes more collateral than the total borrowed amount.

However, annual yield competition is an inappropriate approach to quantifying the probability of the formation of hubs. The most successful companies do not focus on annual yield; they focus on long-term growth. To show that hubs will be successful, we will show how hubs can improve long-term profitability of certain types of existing crypto businesses.

How Hubs Will Form on the Lightning Network

There is no reason to believe that Lightning Hubs will *not* organize in a manner similar to hubs in other networks. There will be hubs that have superior access to the Lightning Network than other hubs. In short, better access to the network means cheaper fees, which means more users, which means cheaper fees, and the positive feedback loop continues [1]. Consequently, a power-law distribution of payment processing, where only a few hubs serve a majority of the network, is likely to exist in Lightning too. It may be the case that hubs eventually follow the current debit card model and there manifests a proliferation of small, well-branded hubs. However, that future is far away from happening and depends on numerous variables such as the regulatory ease of launching a hub (banking laws and money transmitting laws may apply), how hubs can differentiate their product (region-specific branding, privacy guarantees, etc.), and how much mainstream cares about decentralization. In the near-term, Lightning needs to go from zero to one.

Over the course of this analysis, lower and upper bounds will generated to account for the range of possibilities. The independent variable of our models will be the users acquired by the hub.

Revenue

There is one direct way for hubs to make money (charging fees) and other indirect ways (such as improving other, more profitable parts of their business). We will only discuss the direct revenue from hubs in the revenue section. The indirect revenue is reflected later on in this analysis through the incorporation of customer acquisition costs (CAC), as certain types of existing businesses (namely crypto exchanges) will likely launch hubs as a means to acquire more users.

Revenue from Fees, Multi-asset Trading:

The Lightning Network relies on hashed timelock contracts (HTLCs) to ensure that users of the network cannot be robbed by malicious parties. In every HTLC, the receiving party has the right to *not* receive the money sent to them. This is not a problem in traditional, singular-asset payments. However, crypto exchanges cannot currently use HTLCs because traders can easily game HTLCs. If the pre-agreed-upon exchange rate moves against the receiving party, that party can refuse to execute the trade. This effectively turns every HTLC-based, multi-asset trade into an inefficient call option. The inefficiency of a Lightning-based call option is derived from reliance on variable Bitcoin block times.

If exchanges can solve the call option problem, they would be able to immediately turn Lightning hubs into profit centers through trading fees alone. However, there hasn't been any progress to improve this feature of HTLCs. For the foreseeable future, revenue from a hub will need to come solely from Bitcoin payments.

Revenue from Fees, Single-asset Payments:

Transaction fee size will likely depend on the transaction size, similar to the current payment processor model (percentage based system). While many Bitcoin bulls are excited by the possibility of micropayments because of single Satoshi fees, the reality is that this may not be practical for the foreseeable future. Alex Bosworth, the lead infrastructure developer for Lightning, recently tweeted, "If you charge low routing fees to 'help' the network yet you do not route many payments, how much are you really helping? Certainly not with many payments. If you charge higher routing fees to 'help' yourself, but do route many payments, your intent doesn't matter, more get helped."

We will assume that the fee rate charged by hubs has a reasonable range of 0.01% to 0.50%. This range balances Lightning's desire to provide cheap payments with the reality that even 0.50% fees are a significant improvement over alternative payment processors (Visa, Paypal, etc), which charge at least 3%.

The median fee percentage on the Bitcoin mainchain is currently 0.093%, but that does not impact our upper bound of 0.5% significantly because:

- The Bitcoin mainchain cannot scale for mass adoption (as shown by the fee prices during the last bull market--fees will rise significantly from what they are today);
- Lightning hubs significantly improves speed (1 second confirmation times versus 1 hour confirmation times), and speed is a huge value-add, especially to merchants; and
- Lightning will only be used if it provides value-add from cheaper fees and/or speed. Thus, large, non-time sensitive transactions may still use the Bitcoin network due to cheaper fees.

Unlike exchange trades (which have higher processing volumes and thus generate higher revenue than payments), payment fees will not make a hub profitable in the near-term. However, we stated earlier that hub do not need to be profit-centers. Hubs need to provide an avenue for long-term profit growth. One of the most important drivers of long-term growth is customer acquisition.

Customer Acquisition Through Hubs

Hubs are positioned to be an excellent onboarding strategy for established crypto exchanges. This is intuitive: hubs drive Bitcoin's usability and consequently its adoption. Increased Bitcoin adoption means increased demand to purchase Bitcoin with other currencies. To buy Bitcoin with other currencies, users need to access an exchange.

To acquire customers for their exchange business, a hub will also likely provide a front-end payment service in the form of a Lightning wallet (which could contain an in-app exchange). As stated before, hubs will likely differentiate themselves on product offering, specifically where they fall on the scale of user privacy versus government compliance, which will depend on the region(s) in which the hub operates.

The exchange that figures out how to brand their Lightning product will achieve a significant advantage over competing exchanges in the regions where crypto is needed most. Most potential crypto users have no need to speculate on altcoins. Most need crypto to exit from unfair fiat systems. A Lightning Hub would provide that solution. The path to billions of users for Coinbase is not by competing with Fidelity.

We will assume that an exchange's current customer acquisition cost (CAC) is $10 for all of their channels, given that $10 is Coinbase's referral bonus to its users.

Crypto exchanges, most of which are well capitalized, will likely be willing to suffer an economic loss of $10 per user that the hub brings onto its platform. We will herein refer to a hub's user count as x. Thus, a hub most incur no economic loss greater than $10x, where x is the number of users acquired.

We will refer to the direct revenue fees generated from running a hub as z*y, where z is the hub's fee rate and y is the annual amount of transaction volume processed by the hub. We already stated our range for the possible transaction fee rate: (0.01%, 0.5%). The x-y relationship will be estimated across a range of (10, 1000), meaning that the average annual volume per user will lie between $10 USD and $1000 USD. This range accounts for the possibility that large transactions may still be cheaper to execute on the mainchain and that a majority of Bitcoin's 'medium of exchange' users may originate from financially poorer parts of the world. We will model our hub's cost across three examples across each range (the lower bound, the upper bound, and a relevant data point in between the two):

y = 10x; y = 100x; y = 1000x
z = 0.01%; ; z = 0.1%; z = 0.5%.

Cost to Run a Hub

The cost to run a hub can be divided into two broad groups: operational and capital. The operational costs at scale become insignificant for large companies (such as successful crypto exchanges), especially when compared to the the cost of capital. Thus, feel free to skip over this lengthy and more technical section (jump to Capital Costs):

Operational Costs

The operational cost of running a hub can be grouped into the following five buckets.

- Computation resources
    - Storage
    - Internet bandwidth
    - Processing power
- Security
- Legal/regulatory
- Cost of opening/closing channels?
- Miscellaneous

Computation Resources

Hubs will need to either assemble their own hardware or use the cloud. The cloud's cost is much easier to estimate than in-house hardware's, and in-house hardware solutions would only be used if they were significantly cheaper than the cloud. Thus, an in-house hardware solution would only drive the lower bound of this estimate further down. We assumed the usage of Amazon Web Services (AWS) given its widespread popularity and competitive pricing.

On AWS, a hub would need to use EC2, Amazon's compute solution. EC2 allows for the processing of Lightning transactions and provides the memory necessary to store the updated states of each payment channel.

Storage Necessary:

Running a full Bitcoin node: Bitcoin's mempool size has ranged from a few hundred thousand bytes to 0.137 gigabytes over the last three years, while Bitcoin's blockchain size continues to increase as expected, currently approaching 200 gigabytes. A Lightning hub would need to store the current mempool so that it can directly search for newly opened payment channels involving the hub, as well as most of the Bitcoin blockchain. The Bitcoin blockchain grows by

about 13 Gb per year, but the hub could use pruning algorithms to keep the block size constant at 200 GB or less.

Opening and managing payment channels: The hub must store all of its currently opened payment channels, whose size will be directly proportional to the number of channels it has opened. It is assumed that most channels will be a two-party multisignature wallet, whose transaction size on the Bitcoin blockchain can be estimated to be around 500 bytes today but should decrease to around 200 bytes once Schnorr signatures are implemented. It is likely that a hub will need to store each channel opening transaction, as well as some metadata coinciding with it (such the user's account name and any relevant KYC data). We will assume a 3x multiple on the 500 bytes figure (to account for storing metadata) to arrive at 1500 bytes per opened channel. Each user only needs to open one channel with a hub.

Secondly, the hub will need to store an HTLC and metadata for each payment made using the hub. We will assume that on average, each channel has 1 HTLC open at any given time. We will also assume that each HTLC will be about 1000 bytes as well and assume a 3x on the 1000 bytes figure multiple (to account for storing metadata) to arrive at 3000 bytes per actively managed channel. Thus, per user, we estimate a total of 4500 bytes of storage per actively managed channel, with one channel per user.

This brings our total storage estimate to: 200 + (4.5*10^-6)x , where x is the number of users (All variables will stay constant throughout the analysis.).

Note that storage costs are usually the smallest cost associated with node operation, and that only a severe underestimation of HTLC size would impact the broader conclusion of the piece. Also note that a hub may already store the full Bitcoin blockchain for other parts of their business, thus incurring only the additional costs of storing the payment channels themselves.

Compute and Bandwidth:

Bandwidth: All data transferred into an EC2 instance is priced at $0 (see this link for EC2 pricing). Data transferred from an EC2 instance costs $0.09 per GB. Each update to the payment channel requires bilateral communication between the user's device and the hub. Thus, assuming each payment channel is updated 5.5x annually (assuming the same velocity per dollar as before) and each payment is routed between 2 parties, this gives us the bandwidth cost of (4*10^-6)*(*5.5*.09*2)*x, where x is the number of users. This cost will be below $1k per year even at scale and is consequently negligible for this analysis.

Processing Power (and looping back Storage): The cost of the EC2 instance is dependent primarily on the required memory (which we already calculated above) and processing power. Processing power is difficult to estimate, so we modeled it unfavorably to the Lightning network with a lower bound of 2 vCPU (0 users), scaling at a linear rate of 1vCPU per 100k users. This results in the processing power necessary as a function of users in vCPU as: 2 + (1*10^-5)x.

Assuming a hub will model from 0 to 10 million users, storage will range from 200 GB (0 users) to 450 GB (10 mil users), while compute will range from 2vCPU (0 users) to 102 vCPU (10 mil users). Depending on the EC2 instance used, total cost would range between $3.024 per hour at 0 users (r5.12xlarge instance) and $13.338 per hour (x1.32xlarge instance) at 10 mil users. Given that there are 8760 hours in a year, this equates to a range of $26,000 to $115,000 per year.

It is likely that the upper bound is a severe over-estimate, as there are alternative solutions, such as using two c5n.18xlarge instances at less than $8 per hour total, that would have also satisfied our compute needs. That said, it's called an upper bound for a reason. Our lower bound is high because of the need to store the entire Bitcoin blockchain. Also note that this range is dependent only on the number of users and not on the dollar equivalent of Bitcoin locked inside the network.

From these data points, we extrapolated a linear equation to estimate the cost of compute: Compute cost (USD) = 26,000 + 0.0089x.

Security

Hubs need to be accessed in real-time, which means their capital needs to be locked in a Bitcoin hot wallet. A hub's capital is unlikely to be spread across numerous wallets as they benefit from liquidity. The cost of securing one hot wallet will likely begin relatively high per user as the hub needs to have a minimum threshold of employees who can sign off on any transaction. However, users need not worry about this because Lightning maintains Bitcoin's premise as a bearer asset. Hubs can drive down the security cost as much as they want through innovation--it's their money on the line, not their users'. Note that denoting cost per user is less meaningful than denoting cost per dollar for security, so we will use initially use cost per dollar locked to define our security cost.

As the number of capital in the hub grows, it makes sense to occasionally increase the difference between the number of signers *needed* to sign a transaction and the number of permitted signer *permitted* to sign a transaction (Large multi-sig schemes will become will become very feasible/cheap using MAST). The discrepancy between signers required and signers permitted will likely begin at 1 and grow incrementally on rare occasions as more keys are added and the probability of more than 1 key being lost increases past a certain threshold.

We will assume that even with 0 dollars, there are two private keys required to sign transactions (the job life-style would be similar to arbitrage traders at big crypto shops like Galaxy Digital) and three private keys that are permitted (This third private key is held by the hub executives.). We will also assume that four security employees will be responsible for these two private keys (two per key), each needing to be paid $100k per year for this service alone. While the risk of

hub employees going haywire is relatively low (given that they would have to flee the country they reside in), they still need to be compensated well enough to mitigate the risk of theft.

50% of the security employees need to collude to rob the hub. The multiple ɓ represents the likelihood of one security employee colluding based on his/her past. ɓ will likely be driven down close to 0 as only candidates with clean records and strong credit will be able to obtain the job (meaning they are unlikely to break the law and aren't in any financial trouble). We will assume a ɓ of 1%. 50% * ɓ represents the total likelihood of theft per dollar locked in the hub.

*Note that the '50%' figure is rounded for simplification purposes. The more precise decibel would be [(the sum of 2n choose n, 2n choose n+1,...up until 2n choose 2n) / (all possible choices out of 2n)] * ~67% (about one third of these combinations would be irrelevant as many of the combinations account for instances where employees managing the same key collude), which rounds to ~50%.*

The equation for security would be in dollars would be: (50%)*(1%)*q + 400k, where q is the amount of dollars placed into the wallet: Security cost (USD) = 0.005q + 400,000.

The growth in the multisig scheme likely follows a stepwise model once the cost reaches a multiple of 200,000. Thus, the next key (and two employees) would be added once the hub reaches 40 mil USD deployed. Makes sense! However, we were interested in modeling hubs' profits and losses variable upon users, not dollars.

From before, y is the annual amount of transaction volume generated by the hub. Assuming that Lightning payments have timelocks that resolve them within one day, q (the average amount of capital deployed into the hub's wallet on average at any given time) is related to y (annual transaction volume) by q = y/365.

Thus, Security cost (USD) is: 0.005(y/365) + 400,000, where y equals the annual transaction throughput of the hub. From before, we will model hub's profitability across y = 10x; y = 100x; and y = 1000x, where x is the total annual number of users of the hub.

Legal and Regulatory (aka Licenses)

Hubs will need to obtain proper licensing from regulators. They may need a money transmitter license and a banking license to allow for compliant functioning of their business. There is currently regulatory gray area around if hubs would need either of these licenses. Compliant hubs will play it on the safe side, while non-compliant hubs won't care.

Money transmitter licensing to operate across all US territories is roughly $1 - 2 mil up front and an additional 100-300k per year for renewals. It is likely that a hub would already have the necessary money transmitter licensing from their previous business endeavors. All compliant US crypto exchanges already have money transmitting licenses.

Opening a bank costs between $500k and $1 mil if set up in the US, but $150k - 250k offshore. It is likely that a hub would set up offshore or would already have the necessary bank licensing from their previous business endeavors.

Know-your-customer (KYC) costs using a third-party service like World Check could cost up to $5k per month, amounting to $60k per annum. It is likely that a hub already uses one of these services too.

Thus, the upper bound would be $3 mil up front and then $600k per year after. However, the lower range would be 0 if the company already has all the licensing and KYC in place. Based on our assumptions of who is incentivized to run a hub (exchanges), we will put these costs at 0. Either the institution will either have these licenses, or they will try to avoid them by going around governments altogether.

Cost of Opening and Closing Channels

The cost of opening and closing payment channels is dependent primarily on three variables:

- transaction fees on the bitcoin mainchain
- the percentage of the cost passed to the user

The cost of opening and closing a channel depends on the current fees needed to execute a Bitcoin mainchain transaction. The average fee is currently $0.25. This fee would rise significantly if the demand for Bitcoin transactions (either through mainchain uptake or Lightning uptake, which causes mainchain uptake) increased. Thus, this cost on the hub today would be $0.25x, where x is the number of users, but could run as high as $30x.

Because this cost does not improve at the margin and is fairly steep at scale, it is likely that the hub will pass the entire cost to the user him/herself. I would suspect Lightning to follow the gaming industry's model of: first pay to play, then freemium, and lastly get paid through incentive-compatible rewards.

Innovation will continue drive down this cost. For example, channel factories, can allow people to open a new channel for no additional cost when closing a current channel.

Miscellaneous

Other costs include UX development, which may cost several hundred thousand dollars per year depending on the dedicated team's size. However, it is likely that the company that launches a hub already has a plethora of employees working on separate projects, of which a Lightning wallet could be a small feature (such as adding it to the Coinbase Pro Wallet). Thus, annual cost likely has a lower bound of $0 and an upper bound $500k (2-4 employees). Based on our assumptions of who is incentivized to run a hub, we will put these costs at 0.

Total Operational Costs

To summarize, the annual cost range from 0 to 10 million users would be:

- Compute resources: 26,000 + 0.0089x.
- Security: 0.005(y/365) + 400,000; where y= 10x; y = 100x; and y = 1000x
- Legal and regulatory: 0 due to incentives.
- Opening and closing channels: 0 due to incentives.
- Miscellaneous: 0 due to incentives.

Capital Costs

The capital cost will either be reflected on the balance sheet at the LIBOR rate (~3%) or the Bitcoin borrowing rate (~12%). Both are reasonable estimates. Many Bitcoin holders hold Bitcoin for the specific reason that no one can borrow it from them, which is why the borrow rate of Bitcoin is as high as it is. We will assume a perfectly capitalist world where the hub would have lent all of their Bitcoin if they did not set up a hub.

From before, the daily average amount of capital locked in a Lightning hub is q. Also from before, q = y/365, where y is the total annual transaction throughput of a hub.

Annual capital cost (USD) = (y/365)*0.12, for y= 10x; y = 100x; and y = 1000x.

Total Cost

Total annual cost (USD) = Total operational cost (USD) + total capital cost (USD)

Total annual cost (USD) = 26,000 + 0.0089x + 0.005(y/365) + 400,000 + (y/365)*0.12,
for y= 10x; y = 100x; and y = 1000x.

Reiteration of Revenue Generated and Willingness to Spend

To reiterate from earlier, revenue generated directly from hub = z*y, for z = 0.5%; z = 0.1%; z = 0.01%.

So Is it worth it?

Under what conditions for z and the x-y relationship is CAC > total annual cost - total annual revenue directly generated by the hub: **10x > 26,000 + 0.0089x + 0.005(y/365) + 400,000 + (y/365)*0.12 - z*y**. In other words, under what conditions is the hub's CAC more expensive than running a hub.

We show in the Graph 1 below that running a Lightning Hub satisfies this equation at the range of 30k to 45k users depending on the hub's transaction fee (z) and annual throughput per user (x-y relationship). We show in Graph 2 that at scale that the net profit from hubs approaches profitability under most assumptions, turning a previously steep user acquisition cost into an eventually standalone profitable enterprise. Note that both graphs reflect cost, and thus a negative y-axis value denotes a profit.

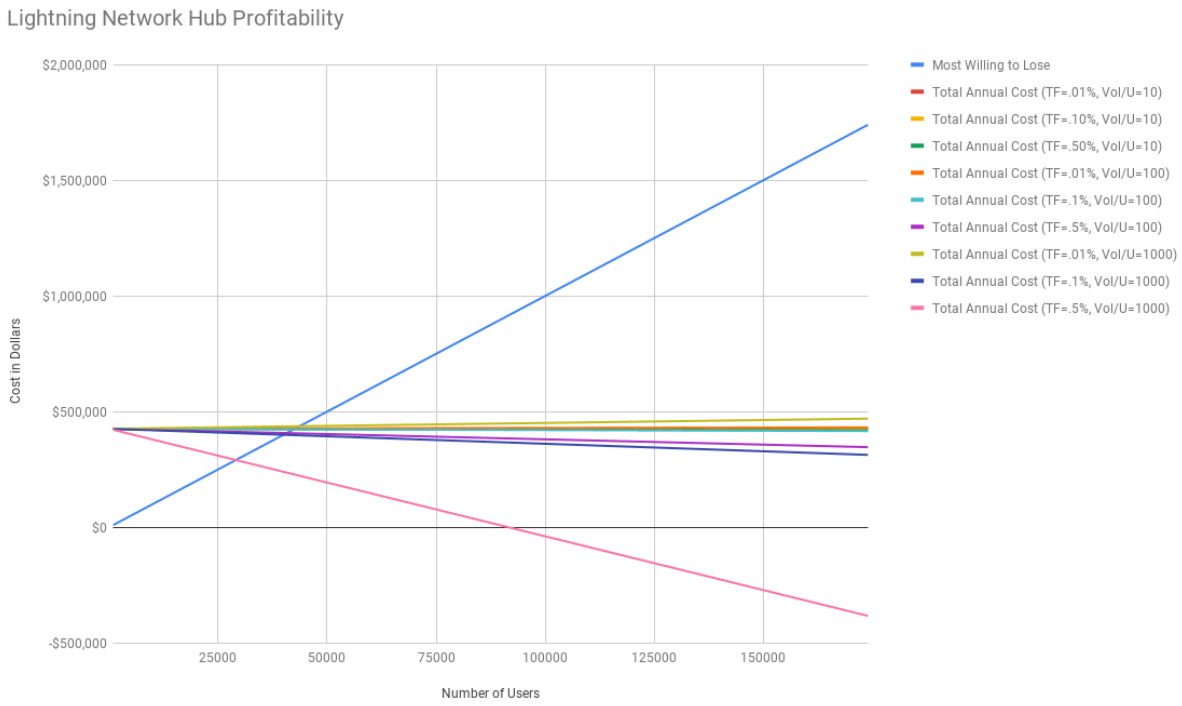

Graph 1

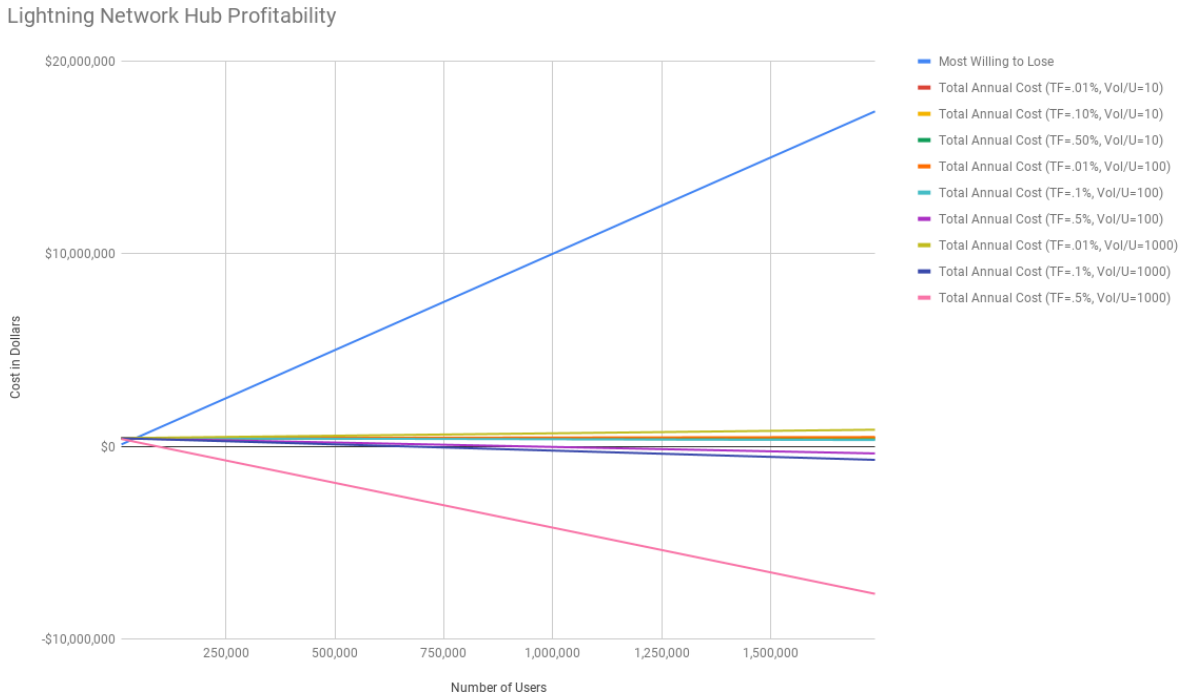

Graph 2

Additionally, our assumptions of writing regulatory costs, channel opening costs, and miscellaneous costs to 0 drastically changed the model. Of these three, the two we may have gotten wrong is our assessment of miscellaneous costs and channel opening costs.

There may need to be additional employees allocated, which has an opportunity cost of labor/salaries. Had those workers/salaries accumulated to $500k USD annually, the breakeven point would have been around 85k users on average. If those costs accumulated to $1mil USD, the breakeven point would have been around 125k users on average.

With regards to channel opening costs, Bitrefill, a crypto payment processor, and Pierre Rochard, a Bitcoin/Lightning developer, currently offer free channels to users that connect with their Lightning nodes. If other hubs offer free channels from the outset, the scale needed to beat current CAC would rise significantly, especially if Bitcoin transaction fees rose. Whether paying users to open channels is scalable depends on the confidential business models of exchanges, which we did not have access to. It may be that our CAC estimate was significantly off, given that payment processor Paypal offered $20 (non-inflation adjusted) for users to join their platform initially.

Hubs could follow the PayPal model, where early users are paid a decreasingly smaller amount as the hub scales. Crypto exchanges already has millions of users, thus their rewards may be smaller than they would have been in their initial stage. However their user growth is still 1000x

from their goal of serving the billions of unbanked, and so they may need to change their rewards program to continue to grow at high rate, which would reflect a new, higher CAC. We estimate that if hubs both pay for a channel and deposit $5 of Bitcoin into a users' account, they will need to scale to millions of users to beat our current presumed CAC. However, a rewards/referral program may allow hubs to scale more quickly.

We are most curious about the highest CAC exchanges are willing to pay. If it is $20 per user, hubs would be able to spend an additional $10 per user based on our current model, which would allow for the opening *and* seeding of user channels. Whether channels are 'pay to play', free, or even pre-funded to incentivize adoption. That offering free services or even paying users from the outset is a even possibility makes us even more confident about the feasibility of Lightning hubs.

Given the current number of active Bitcoin wallets and hubs' potential to significantly increase adoption, we believe that crypto exchanges will launch hubs in an attempt to further improve their businesses. The details of their hub strategy depend on variables that are not made public, such as customer acquisition cost and life-time value of customers. Our research was driven heavily by assumptions due to a lack of public data about exchange businesses and the infancy of the Lightning Network.

Discussion

Security Threats

We would also like to address the belief that the greatest threat to Lightning is its ability to render the Bitcoin mainchain insecure. While the data on how a successful Lightning network would affect miner transaction fees is extremely limited, a highly used Lightning network would likely drive further use of the Bitcoin mainchain from the need to dynamically open and close channels, thus benefiting miner transaction fees. Additionally, it will not make sense for significantly large transactions to be conducted using Lightning, as the fees on Lightning could eclipse Bitcoin mainchain fees. Thus, while potentially a threat, it is more likely that Lightning drives the adoption of the mainchain through the dynamic operation of channels and large transactions.

In the event that Bitcoin mainchain becomes insecure because Lightning hubs are becoming more profitable than mainchain miners, it is likely that hubs would increase their fees to encourage more mainchain use for larger transactions, similar to BTCGuild and Ghash.io's actions when they obtained over 51% of the mainchain's hash rate. When businesses are dependent upon Bitcoin, they have shown to be willing to adjust their profit models to ensure that their profit models remain relevant. An attacked Bitcoin mainchain means an attacked Lightning hub.

What happens to Bitcoin if Lightning Fails?

If Lightning hubs do not exist and Lightning consequently fails, Bitcoin is still significantly more likely than other cryptocurrencies to succeed in becoming the world's non-government money because of its already achieved network effects among users/speculators and ability to copy other projects' features in the open-source world. The greatest threat to Bitcoin as a store of value is Ethereum because of its potential to provide greater utility than Bitcoin by incorporating smart contract functionality.

Ethereum is focused on driving utility to its network to by securing the state and compute of decentralized applications, with the the hope that providing more utility than Bitcoin will ultimately lead it to becoming a better store of value than Bitcoin. The logic follows: the more value that is exchanged using the Ethereum network, the more secure the public chain must be, and security in Ethereum's Proof of Stake system is increased by an influx of dollars (or Bitcoin) into Ether to be bonded on the Ethereum chain. While a very interesting roadmap, Ethereum is less likely to succeed at becoming the global store of value because it has well-capitalized competition for its use case (EOS, Zilliqa, Aelf, Dfinity, Polkadot, Cosmos, etc.) and much of its future roadmap is more complex than Bitcoin's and significantly less tested than Bitcoin's. Complexity creates insecurity; security can only guaranteed after a system is time-tested. It is much more useful for a store of value to be secure than useful. However, I would not count Ethereum out by any means as maintaining its position as the dominant smart contract platform.

Generalized Mining Comparison

Web 3.0 coins hinted at the answer to our initial question of the feasibility of a Lighting Hub's existence. Some of the most well known crypto funds partake in generalized mining (such as Coinfund, Fabric, Notation, Multicoin, etc.). In short, generalized mining is the act of providing the supply-side service for a decentralized network (a hub is the supply-side service in Lightning's case). In most instances of generalized mining, the demand for the service provisioned by the decentralized network approaches zero (such as in the Livepeer network).

However, crypto funds continue to supply services with the goal of kickstarting the network effects needed for service's demand to rise. The cost of capital seldom comes up in these discussions. The annual rate for lending illiquid Web 3.0 tokens to shorters would be much, much higher than the borrowing rate for Bitcoin, especially Web 3.0 tokens with staking models that significantly favor early adopters through high initial inflation rates. **However, it quite clearly becomes obvious that it is rational for crypto funds to operate on a near-term operational and capital loss because it kickstarts the network effects necessary to make their investment successful over the long-term.**

Bitcoin has already 1000xed. However, its growth represents the underlying user growth that has driven the demand for Lightning. Lightning has demand without supply. Crypto funds are less likely to launch hubs because crypto funds do not offer consumer products and

consequently have no customer acquisition costs. A large consumer-facing crypto company, on the other hand, is in a much better position to do so.

Lastly, Bitcoin believers are also significantly more ethos-driven than crypto funds. This is exemplified perfectly by the Bitcoin Cash fork, which is causing Bitcoin whales to lose millions because of their passion for seeing their vision of Bitcoin come true. However, markets cannot rely on generous whales (although one has already taken action). Generous whales are not sustainable. This post focused on businesses incentivized to create hubs because they deem it to be in their best interest for them to grow by an order of magnitude, likely the only sustainable solution to the Lightning Network's scaling.

USD Stablecoins

USD Stablecoins are inherently tied to the US dollar, and as a result do not compete with Bitcoin to be a non-sovereign store of value. However, the two will compete for the use as a medium of exchange in countries with hyper-inflated fiat currencies, especially because the most well-capitalized exchanges have also created their own USD stablecoins.

USD Stablecoins have the advantage of being significantly less volatile historically. Bitcoin has the advantage of being unseizable *and* creatable through the burning of electricity. This second feature allows people to obtain Bitcoin without going through the KYC process, thus furthering its censorship resistance. The only non-seizable stablecoins will be those that are backed by collateral, but these suffer from numerous inefficiencies that may prevent them from scaling (These types of stablecoins appear to behave more like a lending mechanism than a reserve currency.).

It is possible that both USD reserve-backed stablecoins and Bitcoin earn significant market share globally. However, USD stablecoins will not do well in countries that either have recently faced economic sanctions from the US or do not trust the US more broadly, and so Bitcoin has a larger total addressable market. Ultimately, the winner (if there is one) between the two will likely be as a result of numerous macroeconomic factors. Ray Dalio has predicted that the dollar may lose its status as the global reserve currency and as a result decrease in value by 30%. Bitcoin could also depreciate significantly over the same timeframe. Bitcoin's battle with USD denominated stablecoins will not be decided overnight, but rather more likely be a multi-decade competition. Exchanges (turned hubs) would be smart to offer both services initially.

More America

If American-based exchanges successfully operate global hubs, America can retain its economic dominance by continuing to be the backbone of Web 3.0 finance. This partly hedges America against the event that the dollar does not remain the global reserve currency. Thus, American nationals and regulators should be rooting for the international success of American hubs. However, to do this, America will need to forgo enforcing economic sanctions through

American-based hubs. One of the American government's purposes historically has been to ensure the global interests of corporate America.

What happens when those interests conflict with public policy? There could be a compromise, where the US government uses hubs not for sanctions but solely to acquire data and search for wrongdoing, a similar path that the US government has taken with the internet. Blockchains allow value to flow without friction, just as the internet allowed information to flow without friction.

Infancy

The Lightning Network is still in its infancy. It may take significant time for hubs to see mainstream adoption. Today, would people residing in countries with hyper-inflated fiat currencies switch to Bitcoin if Lightning provided cheaper, faster payments than their legacy systems? There's only one way to find out.

Special thanks to Arjun Balaji, Matteo Leibowitz, Nic Carter, Tony Ling, Alex Min, Gary Thung, Tommy Feng, March Zheng, and Roland Li for their ideas and comments.

Endnotes

[1] Farrell and Newman, *Weaponized Interdepence*, http://henryfarrell.net/wp/wp-content/uploads/2018/11/Weaponized-Interdependence_IS.pdf