THE
BITCOIN
STANDARD
RESEARCH BULLETIN

Dr. Saifedean Ammous

April-May 2019, Volume 2, Issue 3

The Bitcoin Standard as a Scaling Solution

Contents

This month's bulletin will focus on the question of Bitcoin scaling, and in particular, how it can grow as a widespread monetary standard. This paper examines the magnitude of the scaling problem, and the challenges Bitcoin will face on its way to a much larger volume of transactions. Due to its nature as a hard money, I argue it is not realistic to expect it to remain a niche network limited in its adoption by on-chain scaling capacity. Demand for hard money is self-reinforcing and will likely make Bitcoin grow far beyond its on-chain scaling capacity, necessitating off-chain scaling solutions. We examine the trade-offs and risks involved in these solutions, and then discuss what can be learned from the growth of the gold standard, and whether Bitcoin can avoid the fate of gold.

Before delving into this topic, a little disclaimer is in order. None of the analysis presented here is intended as a design proposal, and this paper isn't arguing for adoption of one technology over another. Instead, this paper presents my understanding of how economic reality is unfolding around Bitcoin given the technical limitations involved. It is my hope to offer the reader useful mental models with which to understand Bitcoin's scaling and to address the main questions around it, rather than advocating for any particular path for bitcoin users to adopt.

I. The magnitude of the problem

According to The World Payment Report 2018 from Capgemini and BNP Paribas, 482.6 billion non-cash transactions took place around the world in 2016 (about 1.32 billion transactions per day). At a predicted compound annual growth rate of 12.7%, this number is expected to reach 690.8 billion non-cash transactions in 2019, or 1.89 billion daily transactions. For simplicity and a nice round number, let's extrapolate a year further and say we'll have 2 billion daily transactions by 2020.

For comparison, the highest daily transaction volume that the Bitcoin network has ever achieved is 490,459, which happened on December 14, 2017. Currently daily transaction count is around 400,000 transactions a day. At current levels of demand and security, Bitcoin can provide on-chain payments equal to 0.026% of global non-cash transactions. Put differently, if Bitcoin is to handle all global digital payments (at 2020 volumes), it needs to increase its transaction capacity by around 5,000 multiples by next year.

Current bitcoin transaction capacity is being achieved at a block size of around 1 megabyte. The naively obvious approach to scaling simply suggests an increase in the size of blocks until they are large enough to accommodate whatever number of transactions is needed for Bitcoin to take over the world. This was the scaling approach favored by the doomed hard fork attempts Bitcoin XT, Bitcoin Classic, Bitcoin Unlimited and Segwit2x. It was also the driver of the doomed Bcash hard fork (as well as its own even more doomed hard fork, BcashSV). The sorry history of all these poorly thought-out attempts is well worth revisiting in-depth, and Kyle Torpey has written many **articles** on their failures. The important conclusion from all these episodes is that increasing the block size is not a workable scaling solution because even relatively small increases wouldn't move the needle, and would come at the expense of a significant increase in the cost of running a bitcoin full node; this would likely reduce the number of full nodes, which is ultimately the only guarantee of Bitcoin decentralization and immutability.

Bitcoin's core value proposition is its immutability enforced by the consensus rules that full nodes run, which ensures its uncensorable nature and hard monetary policy. A block size increase approach to scaling has proved highly unpopular with bitcoiners, and anyone who attempts it will likely end up with a pointless altcoin like the many dozens of worthless bitcoin forks out there. And even if bitcoiners were to adopt much larger blocks, it wouldn't provide the orders of magnitude increase in scalability needed to for Bitcoin to handle all global transactions.

To handle all global transactions, Bitcoin would need to scale to blocks of around 5 gigabytes each, meaning every computer on the Bitcoin network would need to download this much data roughly every ten minutes (and have the hard drive to store all of these massive blocks), which would accumulate at a rate of almost 0.7 Terabyte per day, indefinitely. This is roughly equivalent to the total hard disk space on today's average commercial

computer, implying that no commercial computer owners would be able to download the Bitcoin blockchain; only people who could afford highly advanced computers would be capable of running a full node. Such a form of Bitcoin would fail to have a large number of people running full nodes, and as a result it would be under serious threat of capture or centralization. When there are only a few dozen full nodes worldwide, it's relatively straightforward to compromise them directly, or to influence them to change the rules of consensus.

Fortunately, other solutions exist that can increase on-chain transaction capacity while avoiding a blocksize increase. Many of the recent improvement proposals promise more efficient transaction handling. But even with all of these improvements, there are hard limits to how many transactions Bitcoin's ledger can record. No matter what optimizations are performed, the bare minimum needed for a single payment to take place is the data needed for the transaction output, which is still 34 bytes of data per transaction. Assuming 4 Megabyte blocks, even the most theoretically efficient use of block space would translate to around 17 million daily transactions, still a far shout from what would be needed for handling all global transactions.

Since Bitcoin's decentralization is the only thing that makes it valuable, its transaction capacity cannot possibly come at the expense of a reduced number of full nodes. Does this mean that Bitcoin is doomed to never scale and remain a niche network processing a few million transactions a day? I would suggest that this is a highly unlikely fate for bitcoin, because hard money cannot stay niche.

II. Hard money cannot stay niche

There is a school of thought that argues Bitcoin must remain a niche and fringe payment network accessible globally, drawing inspiration from Esperanto as a niche global language. Their claim is that Bitcoin will not scale to become a global money given its capacity limitations and government opposition to it. It will only remain useful for people looking to escape capital controls or inflation, and won't ever grow to widespread adoption.

The first problem with this view is that hard money is by its very nature a viral and all-conquering technology that cannot be restricted or restrained from growing. As the first four chapters of my book explain, monetary history is but the history of harder moneys destroying the value of easier moneys and replacing them. A hard money cannot coexist peacefully with easier moneys around it. That state of affairs in itself is an unstable equilibrium that contains the dynamics to alter it. When Europeans

found that west Africans were using beads as money, they took advantage of the fact that the beads are cheap to produce in Europe but expensive to produce in Africa, and brought very large quantities with them to purchase everything valuable in west Africa. There was no way for beads to remain as money in Africa, no matter what the feelings of their holders. Anybody who chose to continue using them as money completely lost their purchasing power; in effect, the beads ceased functioning as a money. The existence of a harder money and other human beings acting in their own self-interest will very severely limit your choice as to the type of money you can use. This is not just about finding someone willing to accept the money you have; more significantly, it is about the consequence to the money you hold from people able to produce it at a cost lower than its market value. As soon as a harder money is found, that money will store value and resist losing it through inflation due to the difficulty of producing it at a cost lower than its market value. That harder money will retain value better than the easy money over time, as its supply increases by relatively smaller quantities.

As the relative value of the two forms of money begins to change in opposite directions, the harder money's pool of available liquidity increases relative to the easier money's pool; in other words, the probability of wanting to trade with someone who is willing to pay with or accept hard money increases. The appreciation in the value of a money results in an increase in its salability, or the likelihood that an individual will be able to sell it when they need to dispose of it. Salability, as Carl Menger emphasized, is the key property of money. Hardness is key to salability because it constantly serves to increase the relative value of the pool of liquidity available for trade.

This process is of course accelerated when people understand it and rationally choose the hardest money. Over time, as more and more wealth goes toward the harder money, more people will want to use it, and demand for it must increase.

The other important example discussed in The Bitcoin Standard concerns the move from bimetallism to gold. For as long as trade in physical coins was the dominant form of trade, silver retained its monetary role due to its superior salability at small scales. But as technology advanced, new forms of money allowed payment in gold and silver without the need to physically move these metals. Paper notes backed by these metals were the most obvious such invention, and other forms of bank accounts and credit instruments also allowed for payment using gold or silver that laid dormant in vaults.

As gold started to also become liquid at small scales, even through intermediaries, there was little reason left to hold on to silver, and its use as a money began to collapse. Even though payments in gold began to increasingly be processed through banking intermediaries, the liquidity of gold continued to grow, along with its value. Even though small payments could still be made with physical silver

coins without relying on banking intermediaries, the liquidity of silver continued to decline along with its value. Once silver lost its raison d'etre as a method of payment for small transactions, there was no reason for two forms of money to continue existing; everyone who used the less liquid money benefited from switching to the more liquid money (and the sooner they switched, the more they benefited).

The lessons from the collapse of bimetallism are applicable to bitcoin and other digital currencies. As soon as gold was usable for all scales of transactions, silver's fate was sealed. That it could still be used for small transactions was no match for the two inexorable forces against its monetary role: the faster supply increase depreciating its value relative to gold, and gold's larger liquidity pool attracting holders toward it and away from silver. Even though many governments had mandated silver as legal tender, they were helpless to stop it from losing its monetary role by the end of the nineteenth century (in yet another fatal blow to The State Theory of Money). Misguided attempts by governments to prop up the price of silver, such as the Silver Purchase Act in the United States, were futile in preserving silver's monetary role; as the value of the national currencies tied to silver plummeted, countries on a silver standard were impoverished.

By the early twentieth century the world was using gold-backed currency, and the growth of gold's liquidity pool further repeled holders away from silver. Even with silver's legally mandated monetary role, its superiority for in-person exchanges without reliance on intermediaries, a monetary role that had lasted for many millennia, and an enormous amount of liquidity held in it until the late nineteenth century, it was to be demonetized in favor of the harder and more liquid money as soon as technology allowed for it. There was simply no reason to hold a different currency less likely to retain its future value, and the market test determined that people preferred the hardness of gold even despite the reliance on an intermediary issuing banknotes (vs the physical silver coins that did not rely on this trust).

This brings us back to the initial comparison between Bitcoin and the World Payments Report statistics. The 482.6 billion transactions mentioned above were specifically called "non-cash transaction" for a reason: they involve intermediaries processing the payment. While these transactions are mostly digital today, that does not make them categorically similar to bitcoin transactions in economic terms. Even though it is digital, a bitcoin transaction is still a cash payment, because the payment is not the liability of anyone. **Bitcoin is a form of cash because only the bearer is able to dispose of it, and they can do so without the need for the consent or permission of a third party intermediary. Bitcoin as digital cash is more comparable to the physical transfer of physical money, such as in-person cash payments, or movements of gold between gold clearing banks or central banks. It is not**

really comparable to the non-cash payments, even though the two might appear similar because they are both digital. The essential quality of bitcoin is that it is a form of payment free of counterparty risk, not that it is digital. Those who expect Bitcoin to grow by displacing intermediated non-cash payment have completely misunderstood its fundamental nature; fortunately, most of those people are no longer involved in Bitcoin, having moved on to some of its doomed forks. If Bitcoin is to continue to grow, it will grow primarily through an increase in the value of the cash payments, or final settlements, it performs. Individual small payments will be built on top of it through secondary layers, and this process is already unfolding. The similarity of this transformation to that of the gold standard was the inspiration behind the title of my book. The movement of bitcoin on-chain is happening for increasingly higher value transactions, and many more transactions are happening on secondary layers (with both lower security and cost).

III. Bitcoin block space supply

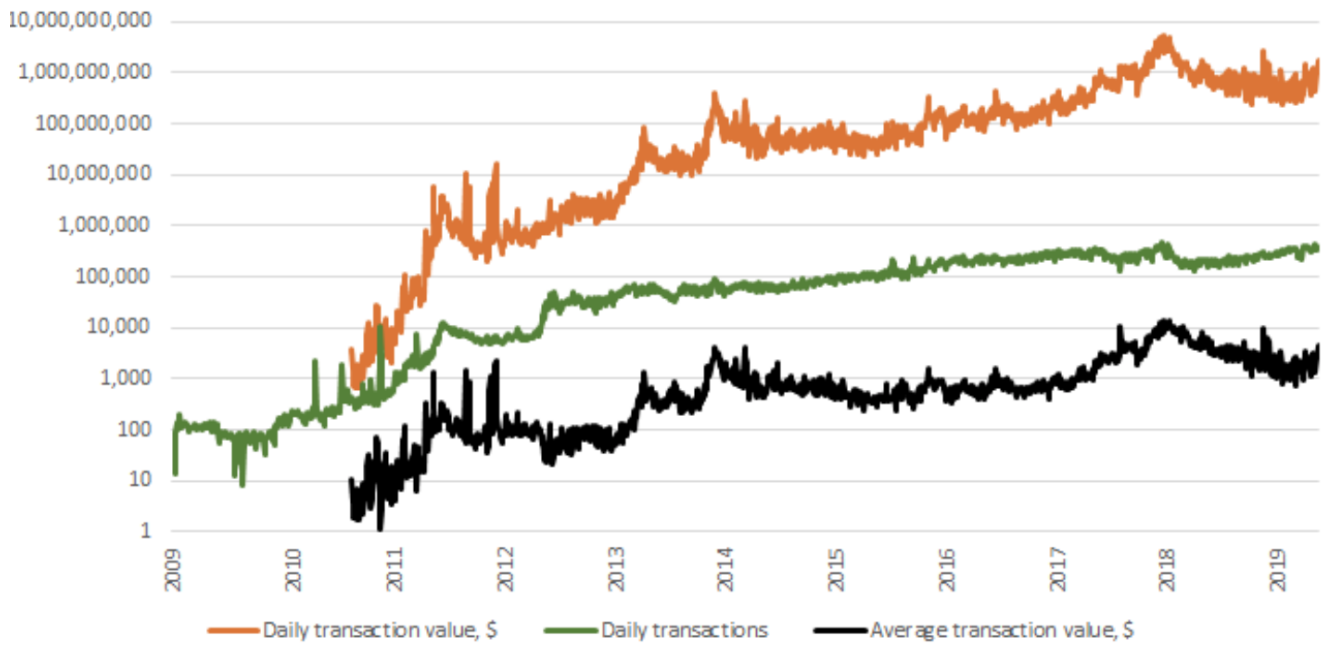

A look at the ten years of Bitcoin's existence shows these trends unmistakably. As the chart below shows, while the number of daily transactions has grown, it is far outpaced by the increase in the value of these transactions. By comparing averages for the most recent three months of data to the first three months of data from 2010 (the first available price data), we find that daily transaction value has increased by roughly 250,000-fold, while transaction count has increased by about 800-fold. **In 2019, the value of the average bitcoin transaction is around 300x the value of the average transaction in 2010.** In fact, the number of daily transactions has largely remained in the range of 200,000 to 400,000 transactions from 2016 until 2019, while the value of transactions has increased roughly 7x over the same period.

As its demand has increased, Bitcoin has not scaled through a larger number of on-chain transactions, but rather by increasing the value of these transactions. Should its demand continue to increase, I expect this trend to continue. With a fixed block size, there is a hard limit to how many transactions can be done on-chain. Even assuming non-contentious forks can increase the block size, they will not be adopted unless they come nowhere near compromising average users' ability to run their own nodes; this means that any block size increase will likely be slow and gradual. Growth in demand for holding bitcoin, on the other hand, does not have the same hard limit. Should bitcoin continue to maintain its core value proposition as a hard money whose supply is perfectly predictable, the growth rate of demand for it will far exceed its ability to handle individual on-chain transactions.

The economics of Bitcoin's block space are a beautiful illustration of market dynamics at work. Its scarce nature necessarily means that a bidding war will ensure only those who value block space the highest will get it. Over time, this pressure has priced out many types of transactions from being registered on-chain, and now more and more are settled off-chain. As was the case with gold and silver, the inability of individuals to use the harder money directly and without intermediaries was not a dealbreaker for them to hold it over the easier money.

Today, many bitcoin-based businesses conduct a majority of their transactions in bitcoin on their own internal databases, and only use the Bitcoin blockchain for final settlement to and from the business. Gambling websites, for instance, will record all bets and winnings on their internal ledgers, and will only use the Bitcoin blockchain when a user deposits or withdraws bitcoin from the website (the same is true for exchanges, where traders speculate on bitcoin and digital currencies). For each on-chain transaction, many thousands of bitcoin-denominated transactions can occur and settle on internal and private ledgers.

This is in contrast to the situation in the earlier days of Bitcoin when betting services such as Satoshi Dice would record thousands of transactions

daily on the Bitcoin blockchain. As transaction fees on the network have risen, these models are no longer sustainable and have changed to rely on the Bitcoin blockchain for settlement only.

Should demand for bitcoin increase significantly, many more uses like this will inevitably be priced out. Because there is no hard limit on its demand, its total daily transaction value can rise to many multiples of today's daily transaction value. If it does, the pool of liquidity for transacting bitcoin will grow, allowing for more valuable purchases and sales to be conducted in bitcoin; this will inevitably price out the transactions of smaller value, as they will not be able to match the transaction fees of these larger transactions.

When considering the types of transactions that will remain on the Bitcoin ledger, it is instructive to think of the alternative avenues available for such transactions. By determining the opportunity cost of not using Bitcoin on-chain for various use cases, we can see which ones can afford to bid the highest for block space. Assuming market participants desire superior security and hard monetary policy, they would be willing to use bitcoin even if transaction fees are significantly higher than alternative payment solutions that rely on trusted third parties and inferior security.

Conversely, if users are not as concerned with superior security and a hard monetary policy for a given use case (e.g. involving smaller value transactions), the opportunity cost of not using Bitcoin is lowered. Currently, individual consumer payments are processed with fees of 0-3% over various payment processors. Given that market participants are less concerned with Bitcoin's value propositions for these use cases, it would only make sense to use bitcoin for these payments if a bitcoin transaction fee were in the cents or at most single digit dollars. Similarly for international remittances, transaction fees are usually tens of dollars, which suggests that as a potential cost ceiling for bitcoin in this use case. If the use of bitcoin for these uses takes off, transaction fees will eventually rise past the cost ceiling, and it would be no longer economical for the users.

This feedback mechanism will continue to price out all manner of uses of Bitcoin's blockchain and will reserve block space only for transactions that need Bitcoin's guarantees the most. As it stands, bitcoin on-chain transactions are a tiny fraction of total bitcoin-denominated transactions, if one were to count trades on exchanges and casinos, as well as all manners of second layer transactions. As bitcoin transaction fees increase, one of the use cases likely to be the most willing to pay will be international final settlement payments between large financial institutions. These are by their nature the most valuable and most security-sensitive transactions today, and the closest thing to a bitcoin transaction currently, in terms of their finality. They currently require days (or even weeks) to complete. Bitcoin is barely beginning to acquire the size and liquidity to allow it to conduct such payments with confidence and security. As it grows it will likely attract

more and more of these transactions, which will crowd out many other use cases and push them off-chain. For some of these crowded out use cases, second-layer solutions will inevitably emerge that retain some of Bitcoin's guarantees while relieving users of its on-chain fees.

IV. Bitcoin cannot replace banks

So far we have examined how the scarcity of the block size necessitates off-chain scaling, but there is another reason why second layer scaling is inevitable: demand for legitimate banking services will continue to exist under a bitcoin standard, just as it has existed under any form of money. Bitcoin block space does not replace the essential functions of banking. There is a lot that is wrong with crony capitalist modern banking, but this is primarily the result of government protection of banks that allows them to profit from unproductive practices and offload the downside risk of their activities to taxpayers.

There are two core functions of banking: holding deposits and allocating investments. The need for these two specialized services is not the result of technical shortcomings of government money that bitcoin could improve upon. They are demanded in a free market for the same reason any good is demanded: consumers value it, and producers specializing in it can provide it at a lower cost and higher quality than individuals could provide it for themselves.

The majority of people with any appreciable liquid savings prefer to have most of their savings deposited with a specialized service that can provide better security. Individuals do not want to have physical possession of their entire life savings at all times because of the risk of loss or theft, and the stress that comes with it. Homes are not designed to optimize for securing large amounts of money, but bank vaults are. It is an inevitable part of human trade and specialization that enterprising individuals would take the initiative and build a facility optimized for securing large amounts of money and employing the kind of security that is unsuitable for a residential home. Individuals would then benefit from paying a small cost to have their money secured at that facility. The benefits are not just that their money is less likely to be stolen, but perhaps more importantly, banking makes it common knowledge in society that anyone with serious money will have it locked up in a safe vault and will not be carrying it on them or storing it at their homes. Even if a criminal were to abduct or extort a millionaire at gunpoint, they would be unlikely to get significant amounts of money from them because most of their wealth would be stored away in a bank, and not immediately available to its owner. This common knowledge is a significant contributor to individual safety, as it severely decreases the likely payoff from violent crime, extortion, and kidnapping.

While bitcoin allows people to send money globally without censorship, it cannot possibly offer them safe and reliable self-custody, as that is inescapably a real world flesh and bone problem. The same censorship-proof nature of bitcoin that allows the sender to irreversibly move money across the world can be utilized by a thief to steal someone's bitcoin. The nodes of the Bitcoin network have no way of distinguishing between different people wielding a private key, and no notion of legitimate or illegitimate ownership of these keys.

No matter the scale at which Bitcoin operates, it is entirely unreasonable to assume it can eliminate the demand of humans to avoid self-custody. That view is built on the naive assumption that people only use banks for payment processing, and so they store their wealth at banks because they need the bank to spend the money. This ignores the demand for storing money for safekeeping and for personal safety.

Importantly, it is also inaccurate to assume that the continued existence of banking will necessarily result in censorship, inflation, and fractional reserve banking. The systematic lack of economic freedom enabled through banking censorship and inflation is the result of the monopolies that governments grant to their banking systems. There is no inherent reason why banking cannot be a normal business where providers strive to please their consumers (think warehouses). Neither is there anything inherently wrong with banking that prevents it from building trusted relationships. People will every day trust strangers to deliver them safe food, drink, and critical tools like cars and airplanes. These industries will function well and consumers will be safe only when a free market exists in these goods, and when consumers have a choice in their providers; this choice forces providers to either care for their clients or suffer the penalty of lost customers and potential failure. As seen in many industries, anytime a government monopoly provides these goods, consumers are in trouble. The problem with banking, then, is not the nature of banking itself, but the fact that it is a government monopoly. In a free market, banking would continue to exist, but would be subject to consumers' choice and their satisfaction. The freedom to choose forces providers to behave their best and swiftly punishes any deviations.

While many bitcoiners themselves have a very strong anti-bank sentiment, and a desire to hold their own money, it does not follow that newcomers entering Bitcoin will necessarily have the same desire or need to follow these ideals. In fact, to impose this model on everyone flies in the face of bitcoin's permissionless nature. Many Bitcoiners may want a world in which everyone gets to be their own bank, but the vast majority of people don't want this anymore than they want to be their own butcher, builder, or baker. There is nothing old-time bitcoiners can do to stop new bitcoiners from banking with bitcoin, if that is what they choose. **It is also inaccurate to assume that the benefits of bitcoin are lost to those who**

choose to deal with custodian services. One may lose the censorship-resistance and permissionless control of owning their own bitcoin private keys, but they nonetheless benefit from holding a hard asset that cannot be inflated away. While there is definitely demand for a permissionless way to send value worldwide, that use case is without a doubt dwarfed by the potentially universal demand for the hardest money. Not everyone has a pressing need for making payments their government does not approve, but everyone will inevitably be compelled by economic reality to converge on the hardest money in the market. As time goes by, and if current trends continue, we can expect demand for holding bitcoin as a hard money to increase even while more transactions are priced off-chain.

The second core function of banking is the allocation of capital through credit and equity investments. The demand for this function is also not something bitcoin can possibly eliminate. The development of banking institutions is an advancement in the process of capital accumulation, allowing for a much more sophisticated division of labor and higher productivity. Because bankers specialize in the deployment of capital, they allow individuals to specialize in their respective fields and focus on being as productive as they can. The individual is freed from the labor of analyzing various investments and assessing their likely returns and risks, since the task is delegated to professionals who specialize in matching individuals' investment goals and risk tolerance with suitable investment projects. The allocation of investment is an act that cannot benefit from the automation and immutability that bitcoin provides to financial transactions. These are activities that require human judgment of factors outside of the Bitcoin blockchain, in relation to subjective individual preferences and desires, and they would exist in any sufficiently advanced capitalist economy.

V. Second layer scaling

Just as transactions with financial instruments based on gold displaced silver coins, it is my contention that second layer bitcoin transactions will in the long run displace transactions that currently take place with easier forms of money. Bitcoin purists may complain that second layer bitcoin transactions will never have the equivalent on-chain transaction security and certainty, but that misses the point. **Second layer bitcoin transactions do not compete with first layer bitcoin transactions, they compete with second layer transactions on inferior moneys.**

The scaling limitations for bitcoin's on-chain volume discussed above make it clear that Bitcoin will probably not scale past a few million on-chain transactions a day, nowhere near the number needed for all individual consumer payments. Bitcoin

itself on its base layer will never be able to handle all of that volume. Further, transactions need about 10 minutes to get a single confirmation on the network, which is highly unsuitable for individuals who expect their consumer payments to complete much more quickly. The level of security and certainty Bitcoin provides for a transaction after it has received a few confirmations is also wasteful overkill for small purchases. For individual small payments, Bitcoin's security is too expensive and wait times are too long. In the same way that payments with gold were standardized and more convenient through banknotes backed by gold, second layer solutions will make bitcoin more predictable, faster, and cheaper, but in the process incur a trade-off of security, liquidity, and censorship-resistance.

While the purists will complain that these kinds of transactions will never have the same level of security as real bitcoin transactions, they cannot do anything to stop the economic reality of individuals preferring these second layer payments with hard money as the base layer to second layer payments on easy money. The limitations that exist will also be present in second layer payment solutions for other types of money. The main difference is that the payment solutions on hard money are likely to allow holders to retain value better into the future. Given the choice between payment solutions on a hard money and payment solutions on an easy money, salability across time dictates that the harder money will inevitably win.

The common mistake that many bitcoiners make when assessing second layer solutions on top of bitcoin is to compare them to bitcoin transactions, but the more accurate comparison is to consumer payment technologies on other forms of money. Conceptually, Bitcoin could scale to handle all of the world's transactions by next week if central banks replaced all their reserves with bitcoin this week. Hypothetically, if the Bitcoin blockchain were only used to settle large transactions between central banks (while they issued currencies fully backed by bitcoin), then all of the world's transactions would effectively be bitcoin second layer transactions. Your government paper money, your checking account, your credit card, and your PayPal account would all become second-layer bitcoin payment solutions in that scenario. Of course, this is not to say that I think such a scenario is likely or even in any way politically feasible; this is just a thought experiment to drive home the parallels between bitcoin and settlement layers.

In the world of national moneys, we cannot really expect central banks to move to bitcoin, as I had discussed in **TBSRB1**. There is nothing wrong with bitcoin that fundamentally prevents its adoption as a reserve asset for central banks. The problem rather lies in the incentives of central banks themselves. I expect that bitcoin is far more likely to grow as an apolitical system independently of the national central banking system, and that this will be healthy for bitcoin in the long run.

As the number of bitcoin holders grows and more people demand payment solutions, there will be an incentive to provide them; the solutions will

be optimized and tailored to work best with bitcoin as it is. This will likely lead to a reinvention of most of the mechanisms we use today for payment. Secondary layer transactions do not share the same level of security as on-chain transactions, but it is not clear why that level of security is needed at all for such transactions. When a customer has an account with an exchange or online casino, they are already trusting that party on many different levels; allowing that party to record transactions on their own ledger, after they've received the deposited customer funds, adds no risk whatsoever. If they choose to exit scam, they could do so regardless of whether their internal transactions were recorded on-chain or off-chain (since funds are only truly under the control of the user after withdrawal from the third party service).

As demand for bitcoin increases, these second layer solutions for scaling will only proliferate, and different levels of risk and safety will emerge for different use cases. Opendimes are another good example. These physical usb keys are made to be tamper-proof, and the bitcoin balance inside them can be verified very quickly. For small sums and transactions between people with a sense of familiarity and trust with one another, this is a very useful mechanism that allows for in-person transactions without needing to be registered on the Bitcoin blockchain. While clearly unsafe for larger sums, it can nonetheless handle a very high number of small transactions and allow for more liquidity in bitcoin transactions.

Multisignature custody solutions will likely also play a role in allowing for cheap second layer payments. Holders could deposit their coins in multisig accounts, such that the coins can only be moved on-chain with both the private keys of the holder and the bank. That bank could then create a payment network for holders of such accounts on its own internal databases to allow individuals to transfer ownership to each other, which would only be settled with on-chain transactions at the end of the day, week, or month.

VI. Lightning

Perhaps the most interesting and promising second layer scaling proposal today is the Lightning Network, which is a new emerging ecosystem of node implementations that allows for an automated, fast, and cheap implementation of a multisig channel-based payment network. Lightning nodes open channels with one another by sending funds to a multisig address using an on-chain transaction. Each party keeps an individual balance on the multisig account, and the parties can pay each other by signing

off-chain lightning transactions that reflect their updated respective balances. When either party chooses to close the channel, an on-chain transaction (reflecting the result of all the off-chain balance updates) is sent from the multisig channel address to the two parties with their respective outstanding balances.
But Lightning users do not necessarily need to build channels with everyone with whom they wish to transact, as payments can be routed through various other nodes and channels to link two parties who do not share a channel. As the number of channels and the liquidity they contain rise, the possibilities of routing payments between users increases. Individual nodes that route payments between nodes can charge routing fees to compensate them for providing the liquidity.

The strength of this approach to scaling is that the setting up and closing of a channel requires just two on-chain transactions in total, and allows both parties to conduct an effectively infinite number of off-chain transactions. Additionally, the timing of the on-chain transactions is flexible, since channels can be opened and closed when demand for on-chain transactions is low. People who establish a pattern of repeated transactions can settle transactions locally on their channel, or through other channels, without having to record every transaction on the Bitcoin blockchain. Despite these benefits, it is important to remember (as Lightning Network engineers such as Alex Bosworth emphasize) that an off-chain transaction on Lightning is not as secure as an on-chain transaction. While most analysis I have seen suggests Lightning is highly secure, it is beyond the scope of my expertise to compare its security to on-chain transactions. I will focus instead on analysing the liquidity of the Lightning Network and how it affects its operation.

The real limitation of the Lightning Network is not in its security or number of transactions, but the depth of the liquidity pool in the network. More people on the network and more money sent to payment channels produce a higher chance that an individual can conduct a trade with someone else on the network (as well as a higher chance that the payment can clear quickly and with low fees). This pool of liquidity, however, is not something that can be solved naturally as the network grows in popularity. The provision of liquidity to the network is a highly complex web of individual economic decisions inextricably linked to people's valuation of time and the inescapable uncertainty of the future.

In page 250 of Human Action, Ludwig von Mises discusses how uncertainty about the future is the key driver of demand for holding money. With no uncertainty of the future, humans could know all their incomes and expenditures ahead of time and plan them optimally to avoid ever having to hold cash. But as uncertainty is an inevitable part of life, people must continue to hold money for future spending.

Committing a balance of bitcoin to a lightning channel is not the equivalent of holding a cash balance, because the money on that channel is

only useful for payment for the counterparty of the channel or others who are connected to them on the Lightning Network, and because establishing channels involves non-negligible costs in fees, time, and coordination. Also, user's channel funds are only liquid to the extent the counterparties in their channel have liquidity. Since liquidity in a channel can generate a return in terms of routing fees, it is more accurate to understand channel balances as an investment to secure routing fees, as well as an option contract: having the right but not the obligation to instantaneously send value through that channel as long as it is open.

Since there is profit to be made from providing liquidity, the optimal liquidity decision for a particular node is not based on individual demand for liquid cash balances, but rather an investment decision based on expected returns from routing fees. If people managed their lightning balances solely based on their need for cash balances, there would be no reason to expect sufficient liquidity to route the payments of others. But since there is a market demand for liquidity, the amount needed to meet that demand will be provided by investment in that liquidity for a return, which implies specialization.

With digital technology, anyone can send a cheap signal to clear a payment and settle it. In reality, the difficult part of payments is the initial deferral of consumption liquidity allows in order to then provide it to those who request it. The job of banks in processing payments can be understood as the provision of liquidity, and in traditional finance they are the ones able to put up cash for payments when needed. As the Lightning Network grows, I believe it will become clear that its growth depends on professional management and provision of liquidity.

The management of the liquidity on channels to optimize for fees is more similar to a specialized commercial enterprise managing liquidity than to individuals managing their expenditure between bank accounts, credit cards, and cash. It is unlikely that an extensive network of liquidity and routing could develop purely from individuals entering into channels with one another, primarily because each individual will be bottlenecked by the liquidity held by their channel counterparties. When an individual opens more channels on the network they create more liquidity for it, but they'll also incur higher costs involved in opening and closing channels. In contrast, opening a channel with a single node specialized in providing liquidity (and with an extensive structure of channels open with many other nodes) will allow that person far more liquidity and reach.

The opportunity to profit from providing reliable liquidity and routing to users suggests that if the Lightning Network were to continue its growth, providing liquidity would likely grow into a profitable and highly sophisticated business. Economic efficiency suggests that the network would be far more robust if liquidity were to become a professional service provided by businesses to consumers.

In such a scenario, one would expect a hub-and-spoke type of arrangement where a global network of nodes with large liquidity all open channels with one another, while individuals would have just a few channels open with these large liquidity nodes. A robust network of nodes each with large liquidity would allow individuals access to cheap and quick routing through deeper liquidity.

Further, if the analysis above with regard to need for custody is accurate, then it is expected that many people will prefer to avoid having to deal with channels themselves, and instead have their bitcoin held in custody by lightning node operators who can also clear payments on-chain.

VII. Trade-offs and risks

The move toward second layer scaling is one that involves risk for users individually, as well as systemic risk for the network. The first and most obvious trade-off is in the network's censorship-resistance. Bitcoin has produced the only reliable technology for transferring value without reliance on intermediaries, and it only manages to do a few hundred thousand of these transactions per day. As demand for bitcoin transaction increases, and individuals resort to second layer solutions that rely on third parties to clear their payments, these parties will be able to censor their transactions and possibly confiscate their coins. One of the main advantages of the Bitcoin network is thus lost for individuals if they choose this type of second layer scaling.

The second risk is more systemic to the network overall, and involves alterations to the network's protocol or consensus parameters. If bitcoin transactions move to second layer solutions where many individuals are trusting third parties to validate their transactions and enforce network consensus rules, Bitcoin deviates from being a peer-to-peer system, and the risk of collusion between nodes processing transactions rises. One can think back to the Segwit2x attempted upgrade and imagine that in a world where far fewer individual users ran their own full nodes, that businesses wanting to change Bitcoin's consensus parameters might have actually gotten away with it had users been reliant on them to enforce consensus rules. If the number of nodes declines, the remaining nodes become more influential and easier to co-opt by attackers or governments. A Bitcoin network with a few hundred nodes is a far less immutable and secure network than one with tens of thousands of nodes.

The risk of losing censorship-resistance is one that each individual needs to assess in contrast to the convenience and cost of other payment and custody options. The other risk is not directly the result of second layer processing itself, but rather a reduction in node count to the extent that jeopardizes

the decentralized nature of Bitcoin. However, the Schelling point of Bitcoin nodes agreeing on the main consensus parameters does not require every user to run their fully-validating node, it merely requires that enough independent full nodes exist to prevent any one particular party from altering the code in a direction it chooses.

What is essential for bitcoin to survive is that the main consensus parameters, particularly the economic parameters, remain immutable, and for that to happen, bitcoin needs a large number of independent nodes unable to coordinate. The larger the number of nodes, the less likely that subgroups will collude. It is not strictly necessary that every individual is able to verify their every transaction on-chain for bitcoin to survive. If the growth of second layer solutions results in a larger liquidity pool for bitcoin, and operating bitcoin full nodes becomes a profitable way to provide banking services, it would financially incentivize the growth of independent nodes, thus making the bitcoin protocol more ossified and harder to change. Not only does the increase in the number of nodes make coordination more difficult, but the profit motive would likely make nodes conservative.

As Bitcoin scales, the challenge will be to introduce second layer solutions that minimize both the trust in third parties and their ability to censor transactions. Yet one must be realistic, and Bitcoin's trustless transactions are not something that can be easily scaled. As discussed above, those priced out of them have no alternatives with the same guarantees. Altcoins have nowhere near the liquidity of bitcoin in the real world, and exist mainly as trading pairs with bitcoin on exchanges. As I have traveled around the world and met many bitcoin brokers, I always ask what percentage of their business is in bitcoin; the answers I have gotten range from 90% to 99% to 99.9%. Altcoins are thus useful for speculating on exchanges, but not so much for the transfer of large sums of money across the world. Most importantly, no altcoin can possibly be viewed as decentralized. Whereas bitcoin is a neutral protocol that only has users, altcoins are subject to small, centrally controlled groups that face no significant barriers to changing consensus rules; this renders altcoins lousy substitutes for Bitcoin's use case as a a long term store of value and a neutral protocol for international payment settlement. Even if an altcoin were to copy Bitcoin's code, it does not follow that it can access the same liquidity and network effects. If bitcoin continues to grow in value due to its hardness as a money, then demand for accessing it as a store of value and for using its large pool of liquidity in trade will mean that second layer solutions on top of it are also likely to be more popular than base layer payment solutions of altcoins. The lesson of the demonetization of silver is again relevant here.

As discussed previously, if demand for bitcoin transactions increases and on-chain transaction fees rise, smaller transactions will be priced out. In my essay from TBSRB3 on the Economics of Mining, I talk about how Bitcoin block space is scarce, and that price is the only way to allocate it. Any approach to

scaling will never alter this fundamental reality. The block space can increase, or it can be used more efficiently, but it will always remain scarce and will never be able to accommodate every transaction in the world. This is the reality of how Bitcoin will operate, and nobody benefits from hand-waving away these trade-offs and limitations, or pretending that they will be solved completely.

If Bitcoin continues to survive and generate increased demand, smaller transactions will find ways of settlement that are not as secure as on-chain transactions. We can imagine a world in which transaction fees continue to rise to the point that only the 1 million transactions that are most willing to pay high fees will be settled on-chain, and everything else will be transacted through less secure means.

Off-chain transactions will never be as safe as on-chain transactions. It took ten years and millions of hours of software development, Satoshi's inimitable genius, and the daily consumption of about as much electricity as Ireland to find a way of doing half a million digital cash transactions daily. We may be able to increase this number marginally over the coming years, but there are no easy ways to increase that number to a global scale, and anyone pretending there are is not being realistic.

The good news is that Bitcoin does not need to be scaled globally on-chain. Bitcoin doesn't have any competitors for trustless, automated, and censorship-resistant global clearance, and the only other asset that comes close to it is gold, whose movement is far more expensive and subject to confiscation, as discussed in **TBSRB1**. Bitcoin needs to be secure and decentralized enough to resist control and capture, and to establish a very clear, broad, and immutable consensus around network rules and money supply considerations. It does not need to accommodate your coffee transactions on-chain.

VIII. Avoiding gold's fate

The ultimate risk involved in scaling Bitcoin is its growing centralization, and history has a great example of that threat: the demise of the gold standard into modern government-controlled debt-based money. The move toward 'second layer' gold payments was its achilles heal in that demise. When payment was predominantly performed in gold and silver coins, individuals had real money in their hands, and governments and banks could not devalue it. As gold-backed notes began to replace physical gold and silver coins for circulation, it became easier for banks and governments to increase the supply of notes beyond their gold backing, which was the root cause for the business cycles (not to mention the unprecedented growth in the size of government and its mandate).

It is too soon to tell if Bitcoin could face a similar fate, but there is one fundamental reason why it has a better chance of resisting this fate than gold: the cost of running a Bitcoin full node is infinitely cheaper than the cost of running a 'gold full node'. I believe the ability of any global settlement system to resist capture or censorship can best be understood as a function of the ease of being able to build an entity that can conduct final settlement with the asset across the world.

The movement of physical gold is expensive, risky, and time-consuming. As trade, technology, and transportation all advanced in the nineteenth century, there was a greater need for transactions, and for netting and settling payments. The inherent nature of the high costs of moving and transporting gold made it inevitable that larger and more centralized banks would have a large advantage over small banks: the larger the bank, the more likely it is to be able to facilitate transactions between two of its clients and avoid the physical movement of gold. Banks also benefited immensely from establishing clearing houses, which later developed into central banks. Eventually the global system of payments around gold was centralized over a few dozen national central banks, each with a de facto monopoly on the clearance of payments into and out of its country. As these reserves became centralized, it became easier for banks to issue fiduciary media unbacked by gold, and for governments to finance themselves through central bank credit. As discussed in **TBSRB1**, the reason fractional reserve banking developed on top of gold was because its expensive clearance gave settlement banks a privileged quasi-monopoly position they could exploit. As reserves were centralized, governments around the world confiscated large amounts of gold from their population by taking over the banking system, where these reserves lay. Since governments co-opted the functions of central banking and forced the acceptance of their currencies, it became effectively illegal to create banking and monetary systems that utilize gold (through many laws and rules discussed in more detail in The Bitcoin Standard). Even today, it is impossible to build a gold-based financial or monetary system, as was seen with the case of e-gold. The cost of setting up a business able to clear gold internationally must also include the steep cost of fighting off the United States government agencies that will attempt to force you to close down one way or another.

The astonishing technical achievement of Bitcoin is that it significantly reduces the cost and duration of international settlement, and more importantly, drastically reduces the set-up cost for performing international clearance. The cost of setting up a bitcoin full node is a few hundred dollars in hardware and bandwidth, and allows anyone to send transactions anywhere in the world. It is even quite possible to perform final clearance with it behind encrypted and undetectable online traffic. Further, recent advancements make it possible to send transactions via radio, satellite, and mesh networks without a regular internet connection. In contrast, any business that wants to clear gold internationally must have a physical location in which the reserves are centralized (and thus subject to government capture).

Transaction costs to settle an international payment using bitcoin is currently far cheaper than using gold. For only a few cents, one can send billions of dollars worth of bitcoin anywhere in the world and achieve final settlement in around an hour or two. And as mentioned above, the more significant advantage bitcoin has is that the cost of setting up a full node is much lower than the cost of setting up a gold clearance bank. Even if bitcoin transaction costs were to rise to many thousands of dollars per transaction, it would still be cheaper and more accessible for millions of people worldwide than the international clearance of gold, which is practically impossible for anyone but central banks and governments.

It is this difference in the initial costs needed to perform international settlement that give Bitcoin the best chance at resisting government control and capture. Whereas with gold we had a system of central banks for settlement, Bitcoin currently has tens of thousands of nodes worldwide each able to perform final settlement. To compromise gold clearance requires a government to infiltrate one building in its territory. To compromise Bitcoin clearance is a herculean task that involves simultaneously quashing many thousands of nodes worldwide that are not very easy to find, but very cheap to replace.

Bitcoin also has the benefit of its public and transparent nature. A bank may be able to easily lie about the amount of reserves it keeps on hand, buts it would have a much harder time credibly claiming ownership of non-existent bitcoin. Individuals depositing bitcoin at financial institutions can keep a close watch on the addresses of that institution, and can publicly audit them. Clients can also place funds in segregated bitcoin addresses under their banks' custody and continue to monitor them. For extra security, multisig solutions can even ensure that the bank is unable to access the funds without consent of the owner.

Acknowledgement

I thank Pierre Rochard for his comments and suggestions on an earlier draft, Adam Tzagournis for his thorough comments, review and editing, and Jamie de Rooij for graphic design.

Thank you very much for subscribing to *The Bitcoin Standard Research Bulletin*.

Please feel free to share this bulletin with any friends you would think might be interested in subscribing to this newsletter, and also, to share excerpts or screenshots from the text on social media.

All the best,
Saifedean Ammous

To subscribe: **www.patreon.com/saifedean**.
Or email **thebitcoinstandard@gmail.com** for instructions on how to subscribe through bitcoin or paypal.