banks without bankers

CRYPTOGRAPHY, INCENTIVES, AND AGENCY IN A BITCOIN FINANCIAL SYSTEM

ERIC YAKES

contact

contact eric at eric@yakes.io

follow him on twitter @ericyakes

get in touch at contact@axiombtc.capital

follow us on twitter @axiombtc

outline

banks without bankers

> *There is a very good reason for Bitcoin-backed banks to exist, issuing their own digital cash currency, redeemable for bitcoins. Bitcoin itself cannot scale to have every single financial transaction in the world be broadcast to everyone and included in the block chain. There needs to be a secondary level of payment systems which is lighter weight and more efficient. Likewise, the time needed for Bitcoin transactions to finalize will be impractical for medium to large value purchases.*
>
> *Bitcoin backed banks will solve these problems. They can work like banks did before nationalization of currency. Different banks can have different policies, some more aggressive, some more conservative. Some would be fractional reserve while others may be 100% Bitcoin backed. Interest rates may vary. Cash from some banks may trade at a discount to that from others.*
>
> *- Hal Finney*

The future of Bitcoin is uncertain. We don't know how well it will scale, how privately it will be used, how it will be stored, or even how it will be used for payments. In addition to the progress of protocols and applications, the development of Bitcoin's financial system may have the most significant impact of all on the value of Bitcoin, the asset. The range of potential outcomes is wide. Consider two hypothetical extremes: In one, all Bitcoin is held in third party custody and users trade receipts between one another. In another, Bitcoin becomes a self-custodial peer-to-peer asset for everybody in the world, providing every conceivable financial function.

Both extremes are unrealistic, and the system will likely end up somewhere in the middle at maturity. Many will pay custodians to store their Bitcoin, and many will not. Some will use protocols they can unilaterally exit, and some will trade claims representing underlying Bitcoin issued by third parties.

Unique to the emergent Bitcoin financial system is the application of cryptography to fundamental financial functions. Novel technologies exist, are being built, and have been theorized that will enable previously unseen functionalities, robustness, and, ultimately, competition amongst Bitcoin financial intermediaries. Key to these novelties is the characteristic of peer-to-peer ("P2P") exchange; that Bitcoin financial intermediaries will surely exist as commercial options, but that newly possible direct manipulation and exchange will exist as well.

I will analyze the possibilities for the development of such a system but will do so with a deliberately partisan lens: I assume as a foundational premise that the greater the P2P possibilities, the better. Better in that I believe financial autonomy is a fundamental good to be strived for; but also better in terms of the stability and neutrality of Bitcoin as whole. It is probably inevitable that trusted third parties will emerge on the basis of convenience provided, but should they dominate their P2P counterparts, the entire system is threatened.

This article is an expansion upon my previous writing Bitcoin Banking (1) which covers the theory behind full reserve banking and free banking, and applies these systems to technologies such as the Lightning Network (LN) and Federated Chaumian Mints. I will expand my analysis of the above, introduce other emerging technologies and focus on the likely economic characteristics of the resulting mix. The best place to start is with a discussion of trust.

if you're going to trust, trust your community...

Few species can cooperate as humans do. We cooperate best with our closest kin as they are most aligned with our genetic interests and genes compete to be propagated into future generations. Evolutionary Biologist John Maynard Smith proposed that genes evolve to find Nash Equilibria when solving strategic problems under competition. Known as the Evolutionarily Stable Strategy, our genes evolve to influence our behavior so that we, generally speaking, help the closest copies of our genes.

Communities in a confined geography also tend to have a relatively greater alignment of interests within than without. As an example, everyone can agree they want security. Debate arises around what methods to use and at what cost.

Genetic alignment varies by location but geographic alignment, by definition, does not. Everywhere in the world, the interests of community members are greatly aligned. There is much to gain from being a part of a community.

As individuals stand to gain more from their communities their risk of loss likewise increases. The Social Risk Hypothesis (2) posits that depression is an adaptive, risk-averse response to the threat of exclusion from social relationships that would have had a critical impact on human survival and reproductive success. It is likely that humans have naturally evolved to avoid social rejection.

There is no denying that people are selfish, and their interests are often not aligned with the interests of their community. All the evolutionary theory in the world doesn't prevent littering. Nor does it prevent the throwing of loud parties at the expense of a neighbor's sleep. And yet, while these examples may create some social friction, they usually aren't deemed costly enough acts to risk social rejection. In contrast, if a community member was caught stealing another's car the social consequences could be much more severe.

Without the cost of community rejection, moral hazard often emerges as the benefits from defection against a conflict of interest outweigh the benefits of maintaining a long-term net-positive contribution. Known as the agency problem, a conflict of interest between a principle and an agent will result in moral hazard, all else equal. Community social costs do not solve the agency problem, but they certainly mitigate it.

Further, communities have evolved with the advent of the Internet. This evolution has rendered geography as a less supremely important characteristic of community alignment while enabling communities with common interests to form globally. Global online communities aren't the result of genetic or geographic alignment. Rather, they form from common interests. The potential for new technologies and financial arrangements to exist among online communities is large, as will be discussed at length below.

Where economic agency exists, community trust can mitigate moral hazard. The advent of the Internet has enabled novel forms of community trust that, in turn, can mitigate novel economic risks.

communities and value

Community trust can be leveraged in a variety of ways. For hundreds (and possibly thousands) of years, informal financial groups have existed as a method for saving and borrowing, be they savings and credit associations, village savings and loan associations,

savings and credit cooperative societies, and so on. Today, informal financial groups are the primary mechanism for savings and borrowing by groups disconnected from formalized financial institutions. (3)

Community trust is also leveraged through formal financial institutions. There are 274 million members of 85k credit unions worldwide as of 2018. (4) Leading up the financial crisis, commercial banks engaged in 5x the amount of subprime lending relative to credit unions and were 2.5x more likely to fail during the crisis. (5) Their public trust is greater and small businesses are 80% less likely to be dissatisfied with a credit union than with a big bank. (6)

According to the FDIC 2020 community banking report, community banks are less likely to close, have performed better since the financial crisis, are a key provider of funding for local businesses (in particular commercial real estate, small business, and agricultural loans), and are more prevalent in rural areas. Community banking is localized by nature. (7)

As brick-and-mortar establishments are uneconomic in many rural environments, digital solutions are being sought to bank the unbanked. Bitcoin is an emergent digital monetary system with properties that can enable the establishment and growth of both informal and formal financial groups. (8) Unique to Bitcoin as a monetary asset is the ability for the individual to maintain self-custody which makes participation in a banking system a choice and not a necessity.

Further, the fact that Bitcoin is digitally native enables voluntary financial groupings to form among the global population connected online. The programmability of Bitcoin enables these groups to innovate novel trust mechanisms. With this technology, community-based financial groups can be formed without geographic constraints. Common interests among geographically dispersed communities can be achieved by leveraging Bitcoin for trade and various financial functions.

The technological properties of Bitcoin enable voluntary adoption among geographically common and distributed communities alike. Novel organizational forms are emerging with the potential to catalyze financial and economic value.

Bitcoin and agency

These emergent systems often require significant user education and specialization. Such burdens are alleviated by entrusting custodial control to service providers – a tradeoff that potentially undermines the systems' purpose. It is natural for economic actors to economize for the sake of specialization. However, if all the Bitcoin in the world was held by 3rd party custodians for the sake of "efficiency," the system would arguably cease to serve its purpose – at the very least, the P2P innovation would be wasted.

Centralized control over Bitcoin custodial operations is a systemic attack vector. In all monetary systems prior to Bitcoin, the transactional efficiencies gained through centralized monetary agency led to moral hazard and ultimately further centralization of the system by political agents. Custodial operations are a step towards centralized agency.

While agency cannot be eliminated today, it can be optimized. The question becomes: who is the ideal agent? One thing all users of Bitcoin have in common is that they are a part of some or other community, and probably of many. A recent realization is that sufficiently knowledgeable trusted community leaders can act as custodians on behalf of the community members as opposed to non-communal 3rd party service providers. For example: a parent who manages a family's finances, the finance department managing a company's expenses, or a group of community leaders managing a community bank.

Bitcoin enables this possibility via the multi-signature transaction – a technology that, in one application, allows for community members to form what is called a "federation". The federated custodial model was theorized by Blockstream (9) and subsequently put into production via Liquid (10) – a multi-signature-based sidechain. The concept behind a federation is simply that multiple participants hold keys that are useless in isolation but can be combined to produce a

signature that is required to make a transaction, thus distributing the trust from one to multiple parties that must cooperate to move funds. (11)

Fedimint, (12) a protocol launched for the purpose of enabling community custody and private transactions, leverages this technology. A primary thesis of Fedimint is that there exists a gap in the market between self-custodial solutions and centralized 3rd party custodial solutions. While many do trust 3rd parties, 2022 was a banner year for demonstrating precisely how this misplaced trust can become horribly consequential. (13) On the other hand, few outside of the community of Bitcoin advocates seem to want to spend the time learning how to self-custody their assets. If we don't trust 3rd parties, but we also don't want to expend the effort to be knowledgeable enough to trust ourselves – what can we do? We can trust our communities.

Trust doesn't scale well, but it can be optimized at the community level. A federation is an enabling technology for those that wish to expend the effort to learn proper custodial practices. They can scale the applicability and utility of this knowledge to the bounds at which trust already exists within their communities. This idea not only fills a gap in the market but has a multitude of implications that could emerge beyond the horizon of localized trust. To understand these implications, first we must understand Fedimint.

Federated custodial technologies leverage cryptography to innovate basic custodial functions. Where agency is necessary, federated custody can exist as a deterrence mechanism against political influence.

fedimint

Fedimint is a protocol at the confluence of 4 primary technologies: (14)

1. **Federations**: groups of individuals that possess computers and provide their memory and processing power to the community. Their computers possess the same software and that enables them to communicate information between one another. The federation is formed by a group of leaders (referred to as "guardians") that generate and control the Bitcoin multi-signature address and also have software that speaks the Fedimint protocol. When users (not guardians... users!) want to join the federation, they are leveraging the federation's memory, processing power, and trustworthiness. This enables them to utilize whatever applications the guardians are providing. Primarily this will be Chaumian eCash (defined just below), but could theoretically be anything, and will probably mostly be financial applications. Federated technology is capable of providing users many things but its primary value proposition is to enable guardians to faithfully execute the protocol on behalf of users.

2. **Multi-signatures (multi-sig):** Bitcoin is stored in a multi-signature address and controlled by the federations guardians. The address requires a threshold level of signatures in order to send Bitcoin transactions. For example, a 3 of 4 multi-sig has 4 possible keys but requires at least three of them to send Bitcoin.

3. **Chaumian eCash:** a private method for representing value that can be traded as a quasi-bearer instrument. (15) It utilizes a cryptographic construct known as a "blind signature": the party issuing (16) the eCash (in this case, the federation) doesn't know the identity of who the eCash is being issued to (the user), yet any third party can nonetheless identify the "signature" on the eCash as having come from this federation. This enables the federation to issue eCash to users that deposit Bitcoin to the federation's multi-signature address. The users hold the eCash on their device (with the ability to hold backups with the federation if they lose their device) making it a kind of trust-dependent digital bearer instrument. There is no public blockchain for the eCash created by guardians. It is simply held in the memory of the user's computer, such as a mobile phone, similar to physical cash, and that can also be backed up to protect against the event of loss. The eCash scheme provides a means of payment that

maintains the censorship resistance of base layer Bitcoin with added privacy but is liable to inflation if a super majority of the fedimint guardians decide to maliciously and covertly increase the supply.

4. **The Lightning Network**: the Lightning network ("LN" hereafter) can ideally be used to forward payments between federations via Lightning gateways (discussed below). This creates the ability to instantly exchange eCash for Bitcoin and has several implications. Importantly, it increases the fungibility between the numerous forms of eCash issued by various federations, reducing the incentive for many to join one federation. Increased fungibility among the eCash of various federations and community trust optimization fundamentally incentivizes systemic decentralization.

The combination of these technologies into a set of rules that the Fedimint software users must follow is what defines the Fedimint protocol. As an open-source protocol, anyone can participate. The ecosystem is comprised of the following participants:

- **Users:** individuals with an app that speaks Fedimint and potentially Bitcoin and LN. They send the Bitcoin to the federation's multi-signature address and receive eCash in exchange. They can send eCash or Lightning to/from any of the applications connected to their wallet; limited only by having the required balance of eCash/ Lightning and if others accept eCash/Lightning.

- **Guardians:** individuals chosen by the community to set up nodes that speak Bitcoin, LN, and Fedimint. They form the federation, manage the hardware, control the Bitcoin in a multi-signature address, and issue the eCash. They can also act as Lightning gateway providers, but this requires specialization (discussed below) and thus another entity called a Lightning service provider ("LSP") will likely fulfill this function.

- **Lightning Gateways**: (17) Lightning node liquidity providers that use Fedimint. The reader can imagine these as a Lightning to eCash exchange that is linked into a fedimint. They integrate with Fedimint users and act as market makers by standing ready to send Lightning payments and receive Lightning payments for a spread. Any federation user can do this, but running a well-connected, high-capacity Lightning node requires specialization and this function will likely be provided by scaled LSPs. If a user wants to "send eCash" to a user at another fedimint, they send the eCash to a gateway, which then forwards along an equivalent Lightning payment to a gateway of the other fedimint, which then sends the receiving user their eCash. eCash can't leave a fedimint – it can only be exchanged for Bitcoin or Bitcoin on Lightning which can be received by gateways to other fedimints and converted back into eCash in the new domain. However, users can integrate with multiple federations and exchange eCash between users of those federations.

- **Modules:** applications within the Fedimint protocol. For a user of a particular federation to use a module, that federation needs to support that module. Fedimints will launch with three standard modules: Bitcoin, (18) eCash, and Lightning Adapter. Examples of potential future modules include smart contract platforms and federated marketplaces. Any federation can choose to support any module. Some federations will have high-performance infrastructure and will support applications that demand it (such as an exchange) while others will have the bare minimum infrastructure that supports the most basic functions of sending eCash and Lightning payments. Users can integrate into however many fedimints they want to utilize and thus, whichever modules they choose.

In summary, guardians form federations that users can choose to be a part of by downloading software that speaks Bitcoin, Lightning, and eCash. The federations with which a user chooses to integrate determine the functionalities they will access. Some will be simple community federations with limited default modules to enable payments. Some federations will have high-powered infrastructure that enables more demanding, potentially commercial scale, applications. Users can custody funds with their community while

linking to commercial-scale federations to use more commercially minded applications. I expect some federations to form among geographic communities and some commercial-scale federations will form to support large scale communities across borders. The system leverages Bitcoin, Lightning, and eCash technology to provide a desirable consumer experience through applications and community custody.

Fedimint is an innovative solution to basic custodial functions. Traditional banking systems have witnessed minimal innovation in custodial operations in recent history, at least functionally. As the most basic function of banking, custodial operations have developed to increase security measures with the proliferation of digital banking. Federated technologies provide a new frontier of innovation for custodial functions. Federated custodial operations have strong potential for growth as well as restructuring the nature of organizations to better align incentives with stakeholder interests. Centralized financial intermediaries must now compete against not only self-custodial systems but also federated systems.

Fedimint combines federated infrastructure with Chaumian eCash, the Lightning Network, and potentially further integrated applications to provide technology that can support all kinds of communities, established or novel.

cashu

Another implementation of eCash is the open-source project Cashu – a non-federated version of Chaumian eCash. (19) Cashu is similar to fedimint in that it issues eCash, yet different in that it is not a federation of servers but rather a single server. While more trust is required without a federation, this system does not require a consensus algorithm which reduces transaction latency. Further Cashu uses only LN, for which no federated approaches exist yet, while fedimint uses both on-chain Bitcoin and LN. Thus, the use cases and demand for Cashu as a protocol are likely to be distinct from those of fedimint.

Notably, Cashu creator Calle (20) has posited a **proof-of-liabilities** scheme (21) theorized to be broadly implemented in eCash systems. Auditing the supply of eCash minted is fundamentally challenging given that eCash ownership, intentionally, is blinded. This topic will be revisited in detail below.

Fedimint and Cashu are both very new, and this discussion is prospective and theoretical for the potential of such an ecosystem. In particular, the integration of LN via LSPs could lay the groundwork for a Bitcoin-native banking system. My first writing on this topic covered the academic theory around this and concluded with a practical discussion. The remainder of this article will expand on this by discussing what could emerge within this ecosystem.

Cashu is a separate eCash protocol optimized for simplicity and speed. Cashu's creator has devised a novel scheme to audit the supply of eCash while still protecting privacy.

monetary utility tradeoffs necessitate various means of payment

Thus far we've defined various protocols that seemingly are implementing forms of money that are distinct from Bitcoin (eCash and LN). In theory, market participants converge upon a monetary standard. In a perfect world, there would only be one form of money. Yet, throughout history this has never been the case.

Why?

While I am not certain that this is conceptually exhaustive, in my book (22) I define three primary reasons for multiple forms of money:

1. **Information opacity**: many different forms of primitive money were used at the same time because neighboring societies weren't economically integrated and were unaware of other forms of money. Awareness is important as it enables individuals to verify the validity of money. As people simply weren't aware of other societies' monies, they weren't able to verify them and would struggle to accept it for trade. As societies have integrated on a global scale, and the Internet has created a global network, the problem of verification has largely been reduced. But not perfectly. Not everyone is connected to the Internet. The level of awareness and ease of verification for a particular form of money is a necessity for widespread adoption.

2. **Sovereign coercion**: today users don't choose money, governments do. If money was chosen in a market and not imposed on society for political purposes, the money of choice would be different to the enforced fiat currencies of today. We are likely witnessing the early stages of the decay of this system, but any transition will require an alternative that is practical enough to use and decentralized enough to eliminate the possibility of coercion.

3. **Monetary utility tradeoffs:** different forms of money maintain different characteristics that make them better for some forms of trade more so than others. For this reason, we often saw dual monetary systems such as cattle and salt or gold and silver throughout history. A contemporary analog could be real estate and dollars where real estate is used to store value while dollars are used for trade.

As a technological innovation, Bitcoin greatly reduces these constraints, but it is arguably not a silver bullet. The Bitcoin base layer network alone (prior to any scaling mechanisms) stores value well but has two primary issues:

1. **Transaction throughput**: the Bitcoin base layer network cannot support global payments as its transaction throughput isn't great enough.

2. **Privacy:** the default setting of Bitcoin is not to be private as transactions are recorded on a public ledger. Significant effort must be expended to increase privacy with Bitcoin transactions. (23)

The Lightning Network is an attempt to solve the transaction throughput problem, although it creates problems of its own. This network is gaining adoption and may become the global payment network required for Bitcoin payments, or at least an important part of such an eventual network. While sending a transaction over LN that is timelocked and fully collateralized in Bitcoin is very similar to sending a direct Bitcoin transaction, it does maintain distinct properties compared to an on-chain Bitcoin transaction. Lightning is faster at the cost of channel capacity constraints required to receive payments. It has weaker security as participating in the network requires storing Bitcoin in a hot wallet, not to mention unknowable protocol risk given Lightning is both newer and arguably more complex than layer-one Bitcoin. To mitigate trust requirements with your channel partner, forcibly closing a channel delays your ability to receive on-chain Bitcoin as well. For these reasons alone, one could argue that the economic properties of a Lightning payment are fundamentally distinct from an on-chain Bitcoin payment and, if one accepts this to be true, it could be argued that Lightning is a monetary medium distinct from Bitcoin.

While theoretically interesting, this may be no more than a semantic distinction. Practically, market participants seem to deem Lightning to be fungible with Bitcoin and this may be all that matters.

Privacy can likewise be addressed in a variety of ways. ECash is one way. It provides nearly perfect privacy but at some cost to auditability. One must trust the issuer of eCash not to debase it (more on this later). However, it truly does provide the anonymity and convenience of physical cash, arguably to an even greater degree as it is digital in nature. For similar theoretical reasons, this could also be defined as a

separate monetary medium – although, again, we will see if this develops any practical relevance. It's important to delineate between a medium of exchange and a means of payment – summarized by Yang (24):

> *The former (medium of exchange) refers to the set of assets in an economy that people regularly exchange for goods and services (a concept of "what"), while the latter (means of payment) is a method that facilitates delivery of money from one to another (a notion of "how"). It suggests that money should be exclusively defined as "medium of exchange," rather than "means of payment." With such a distinction established, one can uniformly explain why currency, demand deposits and smart cards are money (because they are a medium of exchange), and why checks, money orders, or debit and credit cards are not money (because they are only a means of payment but not a medium of exchange).*

Lightning and eCash can also be conceptualized as separate means of payment, rather than as separate monetary media. One could argue that eCash is a distinct asset that derives its value from market participants that demand eCash for its distinct properties. However, its value is ultimately settled on the Bitcoin blockchain. The qualification of eCash as a separate monetary asset or means of payment will be dependent upon how the system exists at maturity. For example, if it were fractionally reserved then its value as an asset will be predicated upon trust in the issuing federation while if it were a full reserve federation then its value predicated upon the purchasing power of Bitcoin. Analogously, US dollars weren't considered gold even when partially backed by gold whereas a 100% reserve gold receipt would be considered closely fungible with owning actual gold (political considerations notwithstanding). Because LN maintains similar economics to owning Bitcoin the asset and seems to be treated as such by users and the market, it likely can be described as a means of payment in Bitcoin.

Theory and semantics aside, the system described thus far would exist at the confluence of three or four protocols: Bitcoin, Lightning, and Fedimint and/or Cashu. The integration of these protocols enables an economy with the security of decentralized Bitcoin as the base layer monetary asset, the privacy and transaction throughput of eCash as a medium of exchange, and unilateral exit from LN channels as a technology facilitating this means of payment.

Various protocols interacting with Bitcoin are forming novel means of payment. Whether or not these ultimately become separate mediums of exchange will be obvious at the systems maturity.

Bitcoin-native money markets

The monetary system described thus far has wide implications for the emergence of digitally native markets. Described in the prior writing, (25) Nik Bhatia theorized that LN is the first Bitcoin-native instantiation of a risk-free rate of interest. While comparable to the reference rate of the fiat system, Lightning is fundamentally distinct in nature as there exists no (economic) counterparty risk by earning yield on Bitcoin through routing fees and liquidity leasing. Bhatia further extrapolates this theory down the risk curve for lending with counterparty risk (see Figure 1):

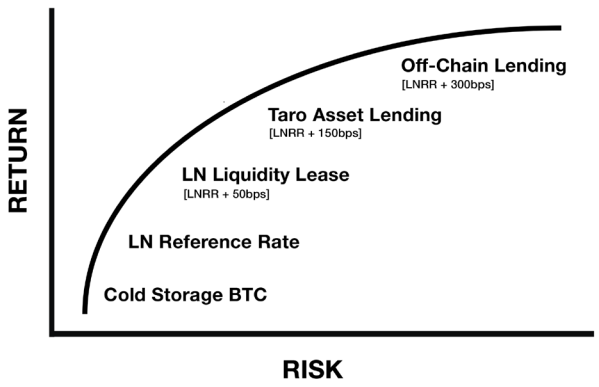

Figure 1 - A novel term structure of interest rates native to the Bitcoin financial system. The Bitcoin Layer (26)

Through this lens, we can view the emergence of LN node operators as the emergence of Bitcoin-native decentralized financial service/infrastructure providers. This will likely be a mix of self-custodial services and custodial services. If custodial service providers evolve to serve banking functions, it could be a mix of full reserve and fractional reserve banks. If LN node operators engage in lending, the market will determine what kind of system ultimately emerges.

What is certain is that money markets are emerging within Bitcoin and market participants are voluntarily participating in them to capture economic benefit. In the US financial system, money markets represent roughly 1/3 of all credit markets by value. (27)

What is a money market? Broadly speaking, money markets are markets that engage in short-term lending of cash. Contrast this with capital markets which are for long term lending, equity investment, and derivatives instruments. Both deal in contracts and the nature of their respective contracts are where the distinction lies (although, again, the distinction is somewhat arbitrary, and we ought not get bogged down in semantics). Capital markets include a greater breadth of assets across a wider variety of contractual terms with a longer timeline. As no non-Bitcoin assets have yet emerged natively within the Bitcoin ecosystem, capital markets have yet to form at scale. However, money markets are forming via the LN.

If federations issuing eCash emerge at scale, then various markets for eCash competing for fungibility with their underlying Bitcoin will also have to exist. The market will determine this fungibility and the primary participants determining it will be Lightning Gateways. They will stand ready to accept eCash and forward an equivalent payment in Lightning to the recipient of a transaction. In doing so, they will discriminate against the various eCash issued by federations. In return for doing so they will earn a spread on each transaction – forming a money market. Thus, a Lightning-to-eCash market maker can earn a spread in return for pricing risk which we can hypothesize depicting on the risk curve (see Figure 2).

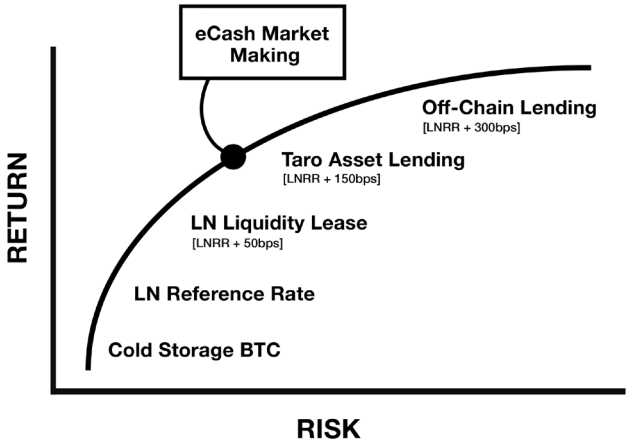

Figure 2 - Market making between lightning gateways and federations can be conceptualized as a new source of economic yield on the term structure of interest rates (The Bitcoin Layer (28) with Eric Yake's additions)

Stated differently, if federated Chaumian eCash finds product-market-fit, the Bitcoin ecosystem will witness a new instantiation of money markets. Markets will emerge trading between Bitcoin or Lightning and the various forms of eCash issued by federations. LSPs can act as brokers – earning a competitive spread between eCash and Lightning market making transactions.

Ultimately the value of these markets will stem from the adoption of the means of transaction they represent. This creates a virtuous cycle of growth. Money markets provide a rate of interest and attract capital. Investment in these markets increases the utility of the functionality they support, which ought to, in turn, increase adoption of the technology.

Bitcoin native money markets are emerging alongside its enabling protocols. Over time these markets will attract investment and create a virtuous cycle of adoption.

the risks of a federated eCash system

ECash is designed to be redeemable for Lightning or Bitcoin via the issuing federation and Lightning gateways are utilized to forward payments between federations, theoretically making the various forms of eCash fungible. The system can be visualized below (see Figure 3):

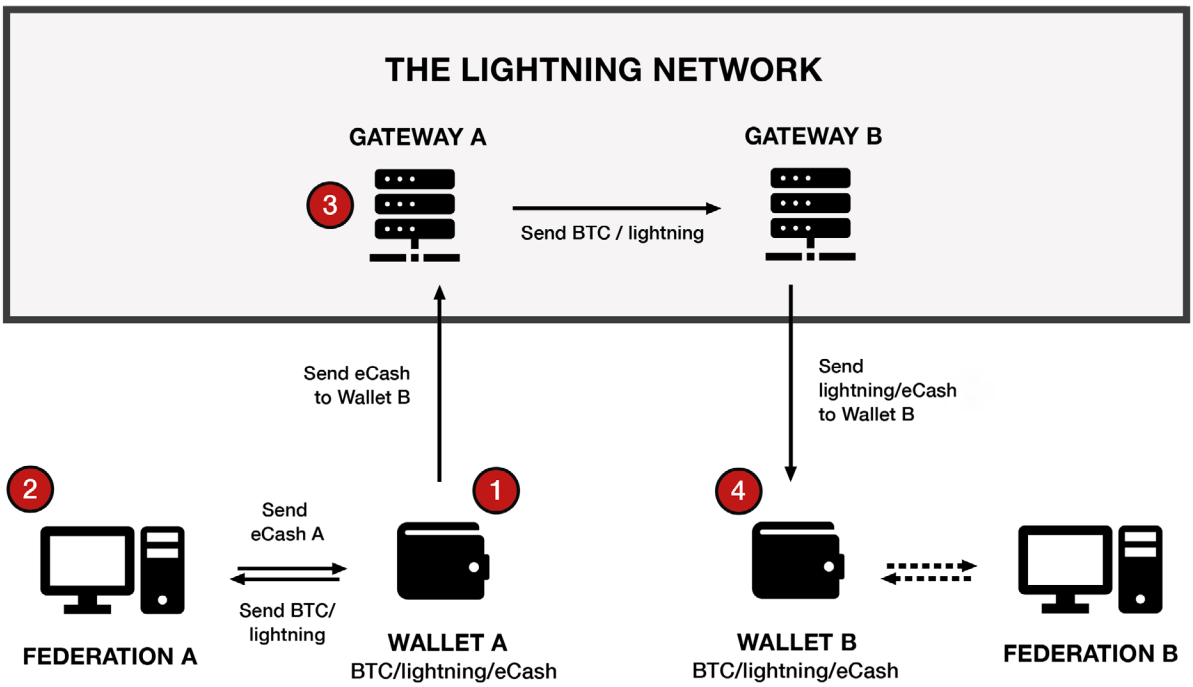

1. Wallet A supports Bitcoin, Lightning, and eCash (of only federation A). It sends BTC to its community's federation.
2. In exchange, the federation sends Wallet A without knowing the owner's identity. Anyone that is a part of federation A can easily receive eCash payments from Wallet A. However, if Wallet A wants to use eCash to send a payment to someone in federation B, they need to utilize a Lightning gateway.
3. The Lightning gateway acts as a market maker, standing ready to send/receive any BTC/Lightning/eCash and earn a spread on each transaction in exchange. So when wallet A sends them eCash, the Lightning Gateway will accept it and forward a payment of BTC/Lightning to another Lightning Gateway connected to Federation B – the federation of Wallet B.
4. Wallet B can then accept this amount in Lightning or eCash and redeem it for BTC at Federation B if desired.

Figure 3 - A simplified visualization of the interactions among user wallets, federated chaumian mints, and LN gateways.

Federations hold an account of BTC and issue eCash. Users of the federation benefit by trusting their custody needs to a specialized manager and being able to transact privately. Self-custody can be complex for the average individual and because Bitcoin maintains settlement finality the risk of losing keys is permanent. Thus, individuals may trade trust in their community federation for the benefits of reducing the risk of lost Bitcoin and gaining privacy in transactions.

However, users aren't just trusting that their respective federation doesn't lose or steal their Bitcoin. Users are also trusting that the federation doesn't issue more eCash than it has received in Bitcoin. Federations can unilaterally issue eCash as there is no cryptographic link to the received Bitcoin. The privacy gains equally imply that the supply is challenging to audit via conventional strategies. Together, this creates the risk of federations debasing the value of eCash. What could prevent this from occurring?

If a community is trusting that the guardians of the federation aren't going to steal their Bitcoin then it follows that they're also trusting the guardians will not debase it. Malicious guardians could simply collude to steal the Bitcoin rather than debase the eCash. However, guardians could potentially use the supposedly trustworthy custodial scheme to slowly debase the eCash (more on this later). That said, there are significant costs to the interests of one's community and these incentives certainly make community custody a less trusted system than 3rd party custody.

On the other hand, what if one's community interests are aligned to debase the value of its eCash? Theoretically, Federation A could gather the community and say that it is going to debase its eCash, send it to Federation B in exchange for goods and services, and distribute the goods received among community members equally. The community agrees because they like the idea of trading nothing for something. However, if this system emerges at scale there will likely be checks and balances that reduce this perverse incentive. To understand this, we can turn to history.

There exists a perverse incentive to debase eCash (in isolation), but natural market incentives mitigate this risk.

Bitcoin and free banking

Free banking systems, discussed at length in the prior writing, (29) can be used as a baseline to assess the competitive dynamics of custodial systems. Applying this understanding to federated eCash systems provides a framework for understanding the potential of this technology.

In free banking, banks can freely issue notes and the market decides if those notes are valuable. If a bank issues too many notes beyond what it holds in reserves, it risks insolvency. Applying this risk to a competitive market limits the degree of note issuance across the system. Circulation credit (30) can only expand so far via note issuance until a systemic bank run is inevitable. However, it was not only in the self-interest of a bank to maintain solvency within the system, but it is also in the interest of the stakeholders of the system. Rational customers wouldn't use a bank if they suspected it of insolvency as this would effectively mean, that as unsecured creditors to an already insolvent institution, they too are insolvent. The problem is that in practice most customers seem to presume solvency and whether or not that is often or ever accurate is another story.

High information asymmetry has historically existed in free banking systems which has led to the collapse of a bank without customers suspecting a problem until it is too late. Because of this, parties who spent the time or naturally had access to more bank-specific information acted as the watchmen of the system. There were three primary groups that limited the degree of note issuance to less than what a bank would naturally issue out of its own perceived self-interest:

1. **Competitors**: competition among banks limits the amount of note expansion one bank can create beyond another. Through the practice of note dueling, more conservative banks would use their

capital to acquire the notes of their competitors that were suspected of high note issuance and subsequently redeem their notes all at once; potentially pushing them into insolvency. The competing bank could then buy the competing institution for cheap and gain market share by acting conservatively. This practice was more common during the infancy of the banking system and declined as the system matured and clearinghouses (discussed below) emerged.

2. **Brokers**: groups that had greater access to bank specific information would speculate on the solvency of banks and profit through arbitrage trading. They would buy notes at a discount that weren't widely accepted and redeem them at the issuing bank for their full redeemable value in gold, profiting from the difference. They could do this because they spent time acquiring specific information about the bank whose notes they intended to broker. This practice created wider note acceptance, placed limitations on the risks that banks could take, and increased information transparency in the system. These broker classes were more prevalent during the emergence of the system. Once the system had reached maturity, clearing houses existed to provide a similar function.

3. **Clearinghouses**: as the system matured, clearing houses emerged to facilitate the functions of brokers and increase information transparency in the system. (31) This constant process of gross note redemption is complicated and operationally intensive, so banks needed a way for netting their redemptions to ultimately reduce the operational burden of the system by settling their debts in one place (or certainly fewer places). This results in the establishment of clearinghouses where all banks go and net their liabilities between one another to settle only the net difference in their accounts. The centralized clearing of debts places the clearinghouses at the center of the system, and they often evolve to serve even more functions such as: credit monitoring, facilitating agreement upon reserve ratios, interest rates, exchange rates, and fee schedules, and assist banks during times of crises (intermediating loans or buyouts). Membership to a clearinghouse was reputation-based and only institutions that met certain standards would be accepted into "the club". This is important as trust is inherent to the system and reputations are paramount to maintaining trust. (32)

Given this, let's return to the problem stated earlier: federations could have an incentive to debase their eCash and trade it for valuable goods and services of another federation. Put simply this is a classic tragedy of the commons, with the commons being trust that the eCash of one federation is fungible with the eCash of another federation. In isolation this incentive appears fatal to the success of the system but when consideration is given to emergent parties, and the checks and balances they apply to the system, natural market dynamics would likely exist to mitigate this risk. Several participants of a federated eCash system like Fedimint could provide these functions:

- **Federations**: most federations will exist simply for custody and payment, but some will exist to provide commercial-scale functionality. We might imagine that we can't have a city where every person has their own road. Custody will eventually emerge as neighborhood streets, city roads, and highways. Fedimint (and LN gateways) provides the architecture and functionality to scale custody into a network of streets and highways. Federations will be competing against one another to garner trust in the broader ecosystem. For streets, it will be a community level trust while for highways it will be a more systemic level of trust and the reputation of a large-scale federation will be paramount to its success.

- **Lightning Gateways**: for a Lightning gateway to integrate with and forward payments of a federation it must hold a balance of that federation's eCash by accepting eCash and forwarding Bitcoin on Lightning to another federation. This will not be an indiscriminate process. Gateways will only act as market makers for various federations if they believe, and can potentially verify, the solvency of that federation. If a gateway notices that eCash balances for it continually increase while

the on-chain data shows the Bitcoin balance has remained relatively flat, they could have a cause for concern. Eliminating their services for a federation could be fatal to the federation's transactional utility. Thus, gateways will only participate with federations whose eCash they feel comfortable holding. Lightning gateways will act as watchmen over the fungibility between eCash issued by various federations out of their own self-interest.

- **eCash Brokers**: it is likely that a broker class will emerge that facilitates a similar function to Lightning gateways but rather than forwarding along Lightning payments they will simply exchange the eCash of federation A for federation B. By acting as direct market makers they would replace the usage of the Lightning network for transaction throughput with a centralized account-based ledger for transaction throughput. Brokers will be constantly monitoring and determining what eCash they want to hold on balance and what they want to either avoid or purchase at a discount. This market making activity will provide another check on eCash fungibility and prevent federations from indiscriminately debasing their eCash value.

- **Proof of Reserves:** companies building technology to monitor the reserves of institutions could also perform a vital function acting effectively as credit monitors of federations. Their emergence can provide certain forms of verification, although not perfect forms. They can certainly monitor the on-chain multi-signature address (the assets) but the liabilities will be more challenging. A federation doesn't know who owns eCash that it has issued but it does know how much it has issued. A federation could provide access and details to 3rd party credit monitors of their history of issuance and redemptions which could provide enough information to assume full reserves or strong solvency (discussed below). Thus, verification is enabled for credit monitoring and reputations of large-scale federations will be paramount towards garnering integration throughout the ecosystem. However, this does not eliminate the risk that a given federation is issuing out of band liabilities which requires 3rd party auditing. For this reason, proof of reserves firms will likely partner with auditing firms or provide services to increase assurances of this risk. Web-of-Stakes (33) is an emergent concept of the Civ Kit protocol that could mitigate this risk in particular applications.

- **Solvency speculators**: a separate class of risk takers akin to hedge funds could emerge that make bets on the solvency of various eCash notes. This would only exist for commercial organizations whereby the fund could execute a redemption attack (34) and hope to profit. This would be similar to note dueling among competitors where the fund benefits not from gaining its competitor's market share but by profiting from a short position on the value of the federation in question. This class would likely be the last to emerge as its existence will be predicated upon established liquid capital markets within the system.

Importantly, the digital nature of this system will enable participants to profit rapidly and cheaply from debasement. By removing the possibility of debasement as a long-term business model, and potentially unprofitable even in the short term, the participants of the system are incentivized to act with prudency. No financial system in history has existed with such an incentive.

If such a system emerges at scale, we'll likely see a consolidation of these functions across various service providers. I anticipate that LSPs could act not only as Lightning gateways but adopt eCash brokerage and potentially acquire or leverage proof of reserve companies and protocols. Just as the brokerage and credit monitoring functions consolidated into clearinghouses of classical free banking systems, so too would I expect consolidation of these functions among community eCash systems. However, all of this assumes that such a system does emerge at scale, which surely will take a long time or not occur at all. Luckily there is potential for technological solutions to emerge and mitigate the risk of eCash debasement in the near term.

Free market incentives align the interests of agents and consumers where trust already exists. This alignment of interests increases as the system matures whereby the value of the system attracts market actors to participate.

proof-of-liabilities scheme for Ecash mints

A federated custodial system (somewhat) mitigates the risk that custodians can steal user funds. It also reduces the risk of the mint debasing the supply of eCash. A free-market system further disincentivizes debasement but for a free market to function most efficiently it requires information to be as transparent as possible. Methodologies that improve the information transparency of outstanding eCash for a mint are paramount for efficient markets. The greater the information transparency of the mint, the greater the auditability of the mint. The tradeoff is that greater auditability can reduce privacy – the purpose of eCash.

Calle, the developer of the Cashu protocol, has proposed (35) a proof-of-liabilities ("PoL") scheme for eCash mints to increase the transparency of eCash supply issuance without, in most cases, reducing the privacy benefits of eCash. (36) This can be achieved by enabling auditability at the systemic level while allowing participants to maintain privacy at the individual level. The system requires three primary voluntary actions of the mint:

1. To publicly commit to rotate its eCash private keys regularly over a predetermined period of time ("epoch"). This allows all eCash in circulation to recycle from old epochs to the current epoch.

2. Produce a publicly auditable list of all issued eCash tokens in the form of mint proofs.

3. Produce a publicly auditable list of all redeemed eCash tokens in the form of burn proofs.

4. A system maintaining these properties can enable users of mints to verifiably detect whether a mint has printed unbacked eCash during a past epoch. It effectively places an expiration date on user eCash and by doing so forces the user to refresh their eCash into the most recent epoch. This expiry of eCash forces users (through automation in their wallet software) to engage in behavior that will ultimately force the mint to report past eCash issuance and redemptions. This is a bit like simulating periodic bank runs at mints. In the words of Calle:

   "In summary, rotating epochs simulates a periodic "bank run" which allows users to observe past epochs and determine whether the mint has manipulated the reports." (37)

Remember, the goal with this scheme is to best ensure that assets (Bitcoin/LN) are greater than or equal to liabilities (eCash) of a particular mint. Placing an expiration on eCash forces the "refresh" of all participant eCash each epoch. So, if a mint is tracking all the eCash it issued, all the eCash it has burned, and is forced to refresh the outstanding amount of eCash each month, then users can publicly verify data on the total supply of eCash that existed during that period of time. The scheme can be visualized in Figure 4 (see next page).

A mint can attempt to cheat in two ways and can be detected in each:

1. Reduce its total eCash issued by publishing as few blind signatures as possible. This can be detected by users when viewing the publicly issued report of blind signatures and noticing that the blind signature of their own eCash is not included. Even a single user can expose a mint for a fraudulent reporting of its eCash issuance. It's important to note that by exposing the mint, the user must forego their privacy guarantees of eCash. (38) However, LN privacy is still strong and even if it wasn't this system, it is still far superior to an account-based ledger system.

2. Increase its total amount of redemptions by creating fake burn proofs. A mint could create a wallet and spend unbacked eCash which it then reports. However, users can prove that a mint is cheating if they can provide a set of tokens whose sum is worth more than the outstanding balance that is reported. This method isn't perfect, and a mint can still theoretically get away with

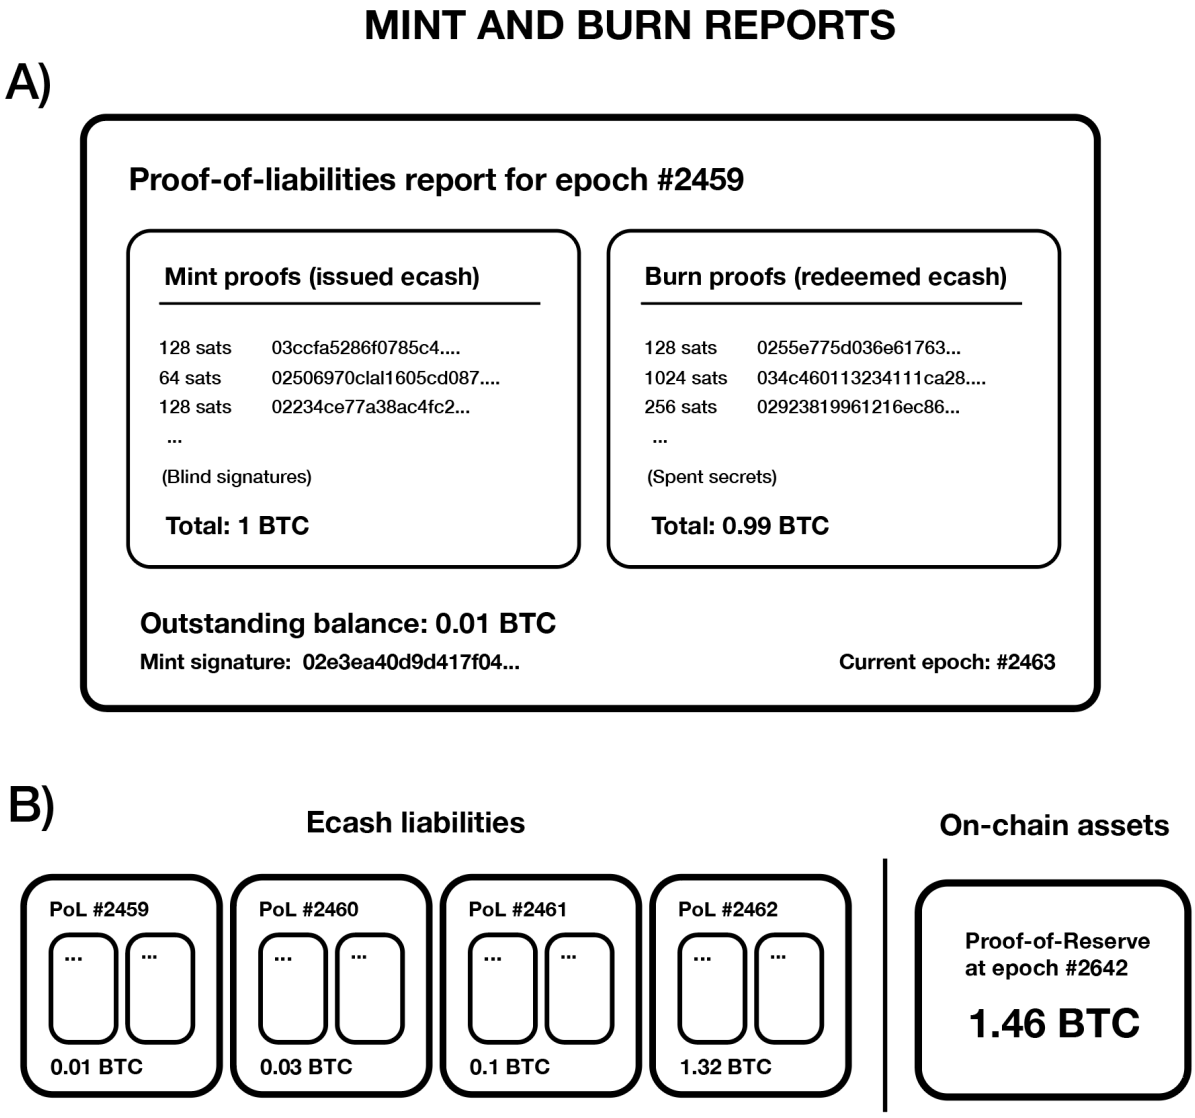

Figure (A) Publicly released PoL reports of the mint include all mint all mint proofs (issued blind signatures) and all burn proofs (redeemed secrets). A Cheating mint would try artifically shorten the list of mint proofs and inflate the list of burn proofs. **(B)** A mint's proof-of-liabilities (outstanding ecash balance) is compared to its proof-of-reserves (on-chain assets). A Cheating mint would try to artifically reduce to open balance but it can't inflate its on-chain assets.

Figure 4 - A proof-of-liabilities scheme for eCash mints for to increase the transparency of eCash supply issuance without, in most cases, reducing the privacy benefits of eCash. Calle (39)

debasement in the short term, but a cheating mint is probabilistically doomed to be caught on a long enough timeline.

Indeed, neither form of eCash debasement auditing is certain to immediately catch a cheating mint. What is certain is that the probability of catching a cheating mint increases as time passes – which is a major innovation. Rational mints are likely to avoid debasement in the first place – knowing the business model is unsustainable without getting caught and risking a run on the mint. Digital note dueling (described earlier) would further exacerbate this phenomenon. The mere knowledge of this risk is likely to act as a deterrence mechanism against any debasement of eCash in the first place.

However, this system requires voluntary action on behalf of the mint as well as voluntary demands of users to participate in mints that maintain such standards. Wallets would need to adopt the necessary technology to enable such a scheme as a best practice. Given that users would need to reduce their privacy to expose a cheating mint, I anticipate consumer protection services (with profit motives) to emerge that are setting up wallets and constantly checking mints for any sort of malicious/negligent action and actively reporting it. (40) Consumer protection agencies would be able to accept the cost of reduced privacy for verification of mint eCash issuance. Rather than mint participants leveraging technology to bear the burden of checking for debasement, a centralized provider could economize on this function and provide a standard check of approval for good-standing mints. For example, proof-of-reserves companies such as Hoseki (41) could evolve to serve such a function. Analogously, if we think of mints as akin to restaurants then there will be Michelin star reviewers that are constantly eating at them without their knowledge and deciding whether or not they get a Michelin star or are reported for poor quality standards.

Innovations from cryptography and clever incentive schemes are removing trust from fundamental economic agency problems.

systemic decentralization deters political capture

Could such a system eventually centralize and be captured by governments just as all banking systems throughout history? I deem it unlikely. If a community custody model emerges it will likely be highly decentralized at the systemic level. If there are 1 billion users and the average federation has 100 members, that would equate to 10 million globally distributed communities. Further, commercial scale federations could exist, for example, as a 70 of 100 multisig, the signers of which are globally geographically diversified. As any participant can join any federation in the world, the competition amongst federations will be severe amongst the most trustworthy federations. Not only will natural decentralization exist for cultural, technical, and geographic reasons, but also because of a high degree of competition. All of this said, the risk of regulatory capture before the system reaches such a scale or has built such dynamics is certainly possible.

The political consequences of clearinghouses and persistence of self-custody at scale will be paramount to a permissionless financial system native to Bitcoin. The ability to practically operate while maintaining Bitcoin self-custody is the primary distinction between a Bitcoin standard system and a gold standard system. Gold wasn't an effective means of payment and thus custodial service, and notes became a necessity, paving the way for fractional reserve banking, and, eventually, the political capture of reserves and their ultimate removal. Bitcoin is different. As more tools emerge for individuals to operate in a self-custodial manner, it will be practical for individuals to conduct

economic activity without entrusting their Bitcoin to custodial providers. It seems likely that this property, unique to Bitcoin, is what will ultimately prevent a central bank- and fiat-like system from emerging.

Lastly, the Fedimint protocol has been designed to fit a particular regulatory niche whereby guardians' custody assets for friends, family, and community interests, where there is no profit motive. If a federation meets these characteristics, it is exempt from financial regulations in many contemporary jurisdictions, but not all. Of course, regulations can change. In the system described thus far, commercial-grade federations could potentially risk regulatory enforcement depending on the jurisdiction. Thus, where a federation exists and what functions it performs will materially influence the financial applications it provides and scale at which it provides them. A positive aspect of this is that regulation will likely act as a decentralizing force among federations.

Federations are a cryptographic and economic innovation that fundamentally incentivize custodial decentralization. Systemic decentralization is paramount to deterring political capture.

the potential of free markets

The system described thus far assumes that in the future eCash becomes fungible (enough) across the federated system for it to function as a widely adopted monetary asset. Why would this occur when Bitcoin and Lightning themselves solve so many problems? I believe that eCash could be value additive to the Bitcoin ecosystem for three primary reasons:

1. **Privacy**: Lightning improves upon the privacy of Bitcoin but eCash provides privacy to a virtually best-in-class degree. Just as the dollar cash system today enables privacy, eCash can do so as true digital cash. Of course, a physical system can be leveraged via private keys printed onto paper with amounts that can be verified via a QR code. Opendime is an example of this and can be traded just like cash. ECash is arguably superior to this physical system in terms of the practicality of the privacy offered in that it maintains optionality for digital payments.

2. **Settlement Finality**: eCash will most commonly reside on one's mobile device and can be backed up with the federation through sharding. The process of sharding is splitting the eCash seed phrase into pieces and sending them to the federation guardians for storage so that in the event of loss, the pieces can be put back together by the guardians and returned. However, this likely would not protect against theft if the thief spends the eCash before the backup process is executed with the guardians.

3. **Capacity Constraints**: the LN is limited by capacity requirements which incentivizes a degree of centralization within the network. ECash does not have this property. If the Lightning Network continues to suffer from inbound capacity limitations, eCash could emerge as a viable payment alternative.

Are these advantages strong enough to justify the potential risk of debasement or even to garner market adoption? Perhaps not. What is important to consider is that this system is not competing against self-custody as much as against centralized 3rd party custodial operations and their adjacent consumer applications, which tend to be much easier to use. The centralized system not only has greater risk of moral hazard but also the risk of debasement. 2022 was a year in which the world realized that many of the assets held on exchanges to which depositors were allegedly entitled simply weren't there. Referred to as "paper Bitcoin," exchanges can be thought of as debasing the supply, but through accounting ledgers that are not fully backed in Bitcoin, and only insofar as these balances are perceived to be real. In any case, the supply of real Bitcoin is unaffected. If centralized ledgers are trusted for 3rd party custody there too will always be the risk

of paper Bitcoin and thus, debasement. Therefore, the question is: would you prefer that the incentive to debase exist in a relatively more decentralized global community model with strong disincentives or through major centralized exchanges with strong incentives for moral hazard?

Even so, will eCash be voluntarily adopted by the market or are Bitcoin and Lightning good enough? Again, perhaps not. Consider that LN is a fundamental component of the federated eCash system. Trading in eCash is optional. Theoretically, one could custody at a federation and simply trade in Lightning. A federation could also be an LSP who could either issue eCash and immediately convert it to Lightning for users or not issue eCash at all. It's possible that a protocol like fedimint most naturally lends itself to a community-based LSP model. However, this would likely result in a much more centralized system as capacity limitations and financial regulations would limit the ability of federations to emerge in small-scale communities. If such a system emerged, it could result in a similar degree of centralization as the 3rd party custodial operations of today. What is important is that the market is able to decipher all of these considerations (and hopefully more that I have yet to comprehend).

ECash is more private and less constrained than the combination of Bitcoin and LN. Further, eCash has stronger protections against debasement than 3rd party custodial providers.

free banking vs. full reserve credit

A particular systemic consideration worth exploring is whether eCash is likely to exist as a full reserve or fractional reserve system, and what would either look like? Consider the below table that expresses the concept of credit by each type and its associated risks (see Table 1):

	ISSUER TYPE		
	Peer-to-Peer (Non-Intermediated)	**Intermediated**	**Private Media Issuer**
TYPE OF CREDIT			
In-kind			
Loan - Secured			
Loan - Unsecured			
Media - Full Reserve			
Media - Fractional Reserve[1]			
RISKS			
Default			
Mismatching - Maturity			
Mismatching - Asset			
Bank Run			
Bank Panic			
Political Capture[2]			

Scale	**Notes**
Does not apply	[1] i.e., Fiduciary media
Applies	[2] Converges toward fiat system
Applies significantly	

Table 1 - An illustrative categorical description of credit types, their primary issuers, and associated risks

The vertical columns show the three primary categories of types of credit issuer: P2P, intermediated, and intermediaries that also issue a form of private media. These three types of issuers can also issue various types of credit and maintain various risks given the type of credit they are issuing. We can see that the issuer type constrains the ability to issue certain types of credit and thus constrains the level of risk associated with the issuer:

1. **P2P**: credit issued in a P2P economy is the most constrained. Individuals can provide in-kind credit whereby no formal loans are given – a service is provided and payment for that service is deferred to a future date. This could be from receiving a beer at your local watering hole that you pay for next week or from accounts receivable contracts that are extended to 60-day terms from 30 days. Loans can also be provided P2P with contracts that secure the loan with collateral or do not. While private credit and loan issuance will surely exist at scale in a P2P economy, lending also necessitates specialization and economization amongst intermediaries. For this reason, intermediation will exist in some form to extend loans. In all of these forms of credit, the issuer is subject to the risk of default.

2. **Intermediated**: credit issued through intermediaries which accept deposits and issue loans. The key distinction in this column is that these intermediaries would be providing loans directly in Bitcoin (or via the LN) and are not issuing their own form of media or means of payment (e.g., eCash). These intermediaries are taking deposits and have terms of their contracts associated. If these are time deposits, then there isn't the risk of a bank run as depositors cannot withdraw until the term of the contract has at least been met. However, if they are demand deposits there is a risk that at any moment depositors can withdraw. Whether or not the intermediary is accepting time deposits or demand deposits, there is always the risk of maturity mismatching – the maturity of the loans the intermediary is making do not match the term of the deposits they are accepting. As long as maturity mismatching exists there is a risk that the intermediary can experience a bank run. Even a full reserve institution with no maturity mismatching could go insolvent from borrowers defaulting on their loans at a loss rate higher than aggregate interest. Further, intermediated systems are more at risk of political capture as they are institutions subject to the laws of various jurisdictions.

3. **Private Media Issuer**: if the intermediary is issuing their own form of private media (e.g., eCash) then they are not making loans directly in Bitcoin (or via the LN) but rather in their particular form of media which is backed with Bitcoin. Intermediaries that issue private media maintain the risk that they manage a fractional reserve. They could also be a full reserve institution that backs their private media 1:1 with Bitcoin. If fully backed, the risk of a bank run is similar to a non-media-issuing intermediary but, if they are fractionally backed, the risk of bank runs is much greater. Further, as history has shown us, fractionally backed institutions are at greater risk of failure which creates a fertile environment for regulatory capture. As regulatory capture increases, the risk of the overarching system eventually becoming a fiat monetary system is greater.

There is potential for all of these systems to emerge within the Bitcoin ecosystem. If we apply this framework directly towards a system of federated eCash issuers we can see that, at scale, the risk of fractional reserve institutions emerging is possible. It is, however, much more dangerous. Efficiencies enabled by the systems digital and cryptographic nature will likely make a fractional reserve a dangerous and unsustainable business model. Further, the bitcoin base supply can't be affected by centralized parties and thus credit issuance cannot be systematically manipulated. Considering these two properties of a bitcoin native financial system, the issuance of fractional reserve credit will be constrained to only localized practices sustainable only in the short term. Fractional reserve is by no means the only way to achieve credit. For reasons discussed earlier, the technical capabilities offered by Bitcoin, Lightning, eCash, and federations are a powerful toolkit to build systems that converge towards a full-reserve standard with strong incentives towards P2P credit.

That said, the emergence of a full-reserve standard isn't a certainty, and a particular incentive could spawn

fractional reserve institutions: seigniorage. Defined as the difference between the cost of issuing money and its value in the marketplace, seigniorage creates a perverse incentive for intermediaries to issue greater amounts of private media on less stringent terms to increase their economic profits. While many aspects of a free banking system limit the degree to which seigniorage can be reasonably extracted, it does not eliminate the incentive to at least attempt it.

However, credit systems can be achieved without a fractional reserve system. The ability to exit the system and operate P2P on Bitcoin and Lightning will be the primary deterrent of unsustainable fractional reserve. As self-custodial P2P economies continue to form in competition with the custodial financial system, it will only become more challenging to competitively operate a fractional reserve organization. Competition with the P2P system is just one deterrent and, along with others described earlier, it has yet to be seen what technologies would ultimately best incentivize P2P or full reserve systems to emerge as the standard.

Topically, the world witnessed how rapidly coordinated bank runs can occur in the online economy of 2023. Information moves at the speed of light and consensus can now quickly form about a given institution's financial health. Mobile banking has made the ease of withdrawals even greater and observably more rapid. Bitcoin is a permissionless asset that can be moved instantaneously via the LN. It is certainly possible that the information transparency enabled by the Internet, mobile technology, the power to unilaterally exit protocols, and Bitcoin will increase the probability of fractional reserve collapsing into insolvency so greatly and over such short time periods that running such an institution will become impossibly risky in practice.

Returning to the concept of community trust, the advent of the Internet has redefined the community to exist not just at a geographic or genetic level but also at a global level defined by a common interest. Bitcoin has enabled online communities as well. Communities are groups of individuals organizing around a common interest and for people to organize they must have the ability to trade. Bitcoin has enabled online communities to permissionlessly organize, trade, and thus form communities. The full extent of this organization and potential for it we have yet to see.

Putting the Fedimint protocol aside, consider the idea that purely the technology of federated custodial models emerges in some form. Such a model necessitates trust and the ease with which it can be applied at a localized level increases the potential for the system to remain decentralized. However, it can also create security in distributed online formats. Various communities that are globally distributed can use this technology to circumvent geographies and form communities in a more secure and less trusted manner.

The incentives are aligned for free banking systems to emerge natively in Bitcoin and technological advancements could prevent fractional reserve banking from becoming a sustainable business practice.

emerging technologies

What technologies could further emerge to enable this theoretical system? Thus far the discussion has covered the LN, federations, and eCash technologies. The combination of these technologies possesses a sufficient range of characteristics to incubate a digitally native financial system, but the system is not complete and could benefit further from emerging technologies.

One technology, currently theoretical in nature, could address several problems within the federation, LN, and eCash system proposed:

- Fedimints don't provide unilateral exit (but Lightning channels do).
- The Lightning Network is structurally geared towards centralization.
- Both fedimints and LN are constrained by block space and fees when onboarding users.

Ark

The emerging protocol, Ark, could address these issues. I'll refrain from explaining the technical mechanics and will focus on the goals of the project, as the protocol is conceptual at this point. (42) Ark uniquely exists at the confluence of multiple other technologies. Like coinjoins, Ark is a mixing service. (43) Like channel factories, Ark is an onboarding mechanism that minimizes on-chain footprint. (44) Similar to how the LN has LSPs, Ark will have ark service providers ("ASPs").

The rough idea is that individuals could onboard to ASPs which allows batches of users to enter the system minimizing their on-chain footprint. This is done by committing bitcoin to a 2-of-2 multisig address with the ASP and receiving a presigned transaction from the ASP providing one of the two signatures necessary to send bitcoin back to yourself. With presigned transactions called VTXOs (Virtual Unspent Transaction Outputs), users of ASPs can swap these with one another for payments. Through this onboarding users can hold their Bitcoin within the ASP and maintain unilateral exit in the event of negligence or malice. This is a solution to the onboarding problem and provides a trustless custodial solution as well.

As a practical example ASPs could be the ideal service provider for acquiring Bitcoin via trustless dollar cost averaging. Imagine herds of thousands of individuals around the world acquiring bitcoin at the same time, on the same scheduling, and all together in the same multisig transaction.

How ASPs will be adopted for payments is uncertain because they require significant capital reserves to support payments. The maximum number of potential payments cannot exceed 10.5M bitcoin because all payment volume transacted through an ASP requires an equivalent number of reserves within a 4-week period. Since there will only ever be 21M bitcoin, at most half of the bitcoin in existence could be used for payments because the other half must be held in reserves to conduct those payments. Depending on the amount of value that 10.5M bitcoin represents will ultimately determine the total transaction volume of the network. The above table expresses a simplified theoretical transaction throughput for the network and the various scenarios that could exist dependent upon the average size of a transaction (see Table 2).

Ark Theoretical Transaction Maximum*	
	Bitcoin
Total Transactions in 4 weeks	105,000,000,000
Average Transaction Value	₿ 0.00010000
Total Payment Volume	₿ 10,500,000
TPS (transactions per second)	***60,764***
Capital Reserves Required	**₿ 10,500,000**
Total Bitcoin	₿ 21,000,000

**Simplified for illustrative purposes*

Table 2 -

An illustrative example Ark transaction (VTXO) economics. This simplified model does not account for total capital reserved within the system costs associated with transaction mechanics.

Transaction Volume Scenarios		
Denomination	Avg. Transaction Size	TPS
Bitcoin	₿ 1.00000000	6
100 Bits	₿ 0.00010000	60,764
Bit	₿ 0.00000100	6,076,389
Satoshi	₿ 0.00000001	607,638,889

The emerging protocol, Ark, could address these issues. I'll refrain from explaining the technical mechanics and will focus on the goals of the project, as the protocol is conceptual at this point. (42) Ark uniquely exists at the confluence of multiple other technologies. Like coinjoins, Ark is a mixing service. (43) Like channel factories, Ark is an onboarding mechanism that minimizes on-chain footprint. (44) Similar to how the LN has LSPs, Ark will have ark service providers ("ASPs").

Putting comparisons aside, what's most interesting about VTXOs as a means of payment is that they become more capital efficient as bitcoin becomes more valuable. There exists some average global payment size which is fixed while the value of bitcoin is expected to appreciate materially (and potentially be ever appreciating). As the value of a bit approaches parity with the global average transaction size, the potential for VTXOs to act as a global payment layer increases. Stated differently, as the bitcoin network represents more value in the world, the ceiling for capital reserves required (10.5M BTC) represents more value, and the average transaction size becomes a smaller proportion of it.

Fundamentally, a key insight here is that the payment networks of tomorrow may not be the payment networks of tomorrow's tomorrow. Payment protocols that aren't constrained by capital reserves might make the most sense today but may not make as much sense as cost of payments declines in proportion to the capital constraints of the system.

That said, the reality of this system will be much more complicated than this theoretical discussion and existing protocols offer distinct and potentially superior payment functionality. The LN can offer less capital-intensive transaction throughput. ECash is the ideal medium for low value high frequency transactions but is trusted and LN is similarly valuable but less private and constrained by liquidity. There are valuable characteristics among all of these protocols that are potentially optimized among federated service providers that speak eCash, LN, and Ark.

Ark, from the perspective discussed thus far, can't be implemented without either the CTV, TXHASH, or elements opcodes soft forks and there are risks to protocol to be deliberated. As discussed, the capital efficiency (and thus capital costs) for payments is a material consideration and potentially a major driver of the protocol's applications. There are other attack vectors to consider as well such as Denial-of-Service (DOS) attacks:

- **Attacks against users:** Exiting an ASP is voluntary but entering an ASP is not. While ASPs provide their users unilateral exit, they also maintain the ability to deny users access or continuous participation because they do not have to onboard a user nor swap payments on their behalf. Similar to a bank, users must trust that they will have access. It is possible that if such a system gained critical mass, blacklists could emerge against participation. Notably, this risk is common to all service providers and development of the P2P system is the solution. Federated infrastructure could be a potential solution to DOS attacks against users – another argument for a multi-protocol service provider optimization.

- **Attacks against ASPs**: DOS attacks are possible against ASPs where arbitrary transactions are conducted to force an ASP to maintain impossibly high or terribly expensive reserve balances. However, the capital costs of such an attack are likely so high that it would only be economic for attackers that maintain a material interest in a competing system to the ASP.

- **Forced Expiration Spam**: (47) stated in the original LN whitepaper (48) and referred to as the Thundering Herd (49) problem where, in LN or any multi-party contractual setup such as Ark, the deliberate or accidental failure of a large, dedicated user requires many other users to put many time-sensitive transactions onchain all at the same time. In the case of Ark, if the ASP goes offline permanently then everyone needs to exit before their funds are forfeited to the ASP. It's possible that this problem could cause severe user and network issues precisely when they are least desired.

That said, today Ark appears to be distinct and viable. The most interesting consideration of Ark is that it

could be the protocol necessary to achieve trustless banking. The lightning network can be utilized without a service provider, but Ark cannot – much like a bank. VTXOs (Ark presigned transactions) are another means of payment but cryptographically guaranteed to be fully reserved. Analogous to a trustless cashier's check, VTXOs are not just as good as gold, they're as good as bitcoin. All things considered; the Ark protocol could provide the necessary infrastructure for a trustless free banking system of service providers to emerge – removing agency from fundamental economic functions. From this, a spectrum of functionality emerges for basic financial functions from fully centralized custodial operations to custodial LSPs to fedimints to trustless ASPs to purely P2P systems.

By applying cryptographic solutions to fundamental financial functions, agency can be reduced throughout the financial system. Emerging technologies native to the Bitcoin ecosystem continue to express novel cryptographic innovations.

final thoughts

Many of the concepts discussed throughout this article are theoretical in nature while some are real technologies solving real problems today. A protocol stack is emerging as the basis of Bitcoin-native banking and general financial services. Part of this stack is layered protocols that maintain unilateral exit while other parts are trusted designs.

Imagine a system where users dollar-cost-average into Bitcoin via Ark, use federated technology for custody, use eCash as the private cash balance for everyday transactions, and on the backend all service providers are clearing balances between one another via the LN. Fedimints and ASPs could act as banking infrastructure and the LN could act as the clearing houses amongst them as a hub and spoke model. Further, competition and information transparency are fostered by technologies like Web-of-Stakes reputation management systems and expiring eCash proof-of-reserve systems.

A P2P self-sovereign financial system is irreducibly complex. Centralized systems are required to scaffold towards decentralized systems. However, with the necessary infrastructure in place, a self-sovereign system could spawn digitally native capital markets at scale and could contribute to entrenching Bitcoin as a standard unit of account. The more settlement that occurs in Bitcoin, or which is linked to Bitcoin in some way or another, the greater its chances of acting as a unit of account.

The described system is strikingly parallel to historical banking systems but leverages the power of cryptography to remove agency from a variety of financial functions and optimize for it where still necessary. The complexity of incentives is a major risk to this vision but if information is transparent and competition is high, free markets will solve this issue in a way that no individual can fathom.

At the end of the day, what matters is that the option of decentralized custody and functionality around P2P finance preempts an otherwise inevitable centralization in banking, and that censorship and inflation alike become obsoleted by technology.

A special thanks to Allen Farrington for his considerable time spent reviewing, advising, and including his original thoughts in this writing.

Thanks to all others who helped review and provided thoughts on many of the addressed concepts including: Obi Nwosu, Calle, Lloyd Fournier, Alex Berge, Matt Black, Alex Lewin, and Burak Keceli

End Notes:

(1) Bitcoin Banking Systems: Full Reserve vs. Free Banking, Eric Yakes. Available: https://yakes.io/Bitcoin-banking-systems-full-reserve-vs-free-banking/

(2) Covered in academic review available at: https://www.sciencedirect.com/science/article/abs/pii/S027858460600008X?via%3Dihub

(3) The New Microfinance Handbook: A Financial Market System Perspective, Chapter 6; Joanna Ledgerwood, Julie Earne, & Candace Nelson. Available: https://openknowledge.worldbank.org/entities/publication/f04e0858-2720-5ccb-a83f-950d215e1bc6

(4) The Competitive Edge of Credit Unions in Costa Rica: From Financial Repression to the Risks of a New Financial Environment; Miguel Rojas, Sébastien Deschênes, Lavasoa Romboarisata, & André Leclerc. Available: https://anserj.ca/index.php/cjnser/article/view/289

(5) "How Did Bank Lending to Small Business in the United States Fare After the Financial Crisis?, Rebel A Cole. Available: https://www.sba.gov/sites/default/files/439-How-Did-Bank-Lending-to-Small-Business-Fare.pdf

(6) Small Business Credit Survey 2016, Federal Reserve Bank of New York. Available: https://www.newyorkfed.org/medialibrary/media/smallbusiness/2016/SBCS-Report-EmployerFirms-2016.pdf#page=23

(7) FDIC 2020 community banking report, FDIC. Available: https://www.fdic.gov/resources/community-banking/report/2020/2020-cbi-study-full.pdf

(8) The Growr protocol is an emerging example of this: https://www.growr.xyz/

(9) Strong Federations: An interoperable Blockchain Solution to Centralized Third Party Risks; Johnny Dilley, Andrew Poelstra, Jonathan Wilkins, Marta Piekarska, Ben Gorlick, and Mark Friedenbach. Available: https://blockstream.com/strong-federations.pdf

(10) Liquid: A Bitcoin Sidechain; Jonas Nick, Andrew Poelstra, Gregory Sanders. Available: https://blockstream.com/assets/downloads/pdf/liquid-whitepaper.pdf

(11) A less technically involved version of multisignature-based custody is Unchained Capital's "collaborative custody" model, which represents a similar set of targeted trade-offs in trust to make custody less risky (and frankly less scary) for individuals.

(12) n.b. I use "Fedimint," upper-case, to refer to the protocol and "fedimint," lower-case, to refer to individual instantiations of the protocol.

(13) As covered extensively in previous Axiom article, Green Eggs And Ham –https://static1.squarespace.com/static/62de2a644f0418669484e364/t/64b2868481cb83591311ad27/1689421454386/green.pdf

(14) Fedimint.org has a more involved primer, available: https://fedimint.org/docs/intro

(15) A mathematical explanation of which can be found in the original whitepaper Untraceable Electronic Cash by Chaum, Fiat, and Naor, available: https://blog.koehntopp.de/uploads/chaum_fiat_naor_ecash.pdf

(16) In this context, the concept of issuance is defined by the mint producing a signature. It is actually the user who "creates" the eCash, but it is not usable writes it's signature upon it. So, the term "mint" is a bit of a misnomer but the signature of the mint is what gives the eCash its value.

(17) The Lightning network (LN) is a secondary payment layer that optimizes around transaction efficiency for payments. Bitcoin owners can lock their Bitcoin into a channel with a channel partner (i.e., another person agreeing to conduct transactions with you) and send payments back and forth with that person in a much more scalable way, not to exceed the amount of Bitcoin you both have committed to the channel. Further, one can use the connections of their channels to forward payments to others that they don't have an active channel with, for a fee. These forwarded payments are routed by the various nodes of the LN until the payment reaches its end destination. Each node forwarding the payment (i.e., routing the payment) receives a small fee for doing so.

(18) Fedimint modules are the current implementation of the Fedimint protocol. Fedi, the company, is building the initial fedimint application utilizing these standard modules as well as adding additional modules of their own.

(19) This project trusts a single counterparty with custody of your Bitcoin and is experimental. Do not send any amount of Bitcoin you aren't comfortable with losing to this mint. You can review the project here: https://cashu.space/

(20) Twitter account available: https://twitter.com/callebtc

(21) GitHub documentation available: https://gist.github.com/callebtc/ed5228d1d8cbaade0104db5d1cf63939#file-ecash-pol-md

(22) The 7th Property: Bitcoin and the Monetary Revolution, Eric Yakes. Available: https://www.amazon.com/7th-Property-Bitcoin-Monetary-Revolution/dp/0578902621

(23) Bitcoin transactions refers to all potential transactions, not simply base layer transactions. The base layer likely doesn't require privacy to function as a base settlement layer but it can only do so if there are other means of payment that provide privacy. Such means of payment are likely enabled througha combination of layers and adjacent protocols.

(24) What is (Not) Money? Medium of Exchange ≠ Means of Payment; Bill Z. Yang. Available: https://journals.sagepub.com/doi/abs/10.1177/0569434507051002l3?journalCode=aexb

(25) Bitcoin Banking Systems: Full Reserve vs Free Banking, Eric Yakes. Available: https://yakes.io/Bitcoin-banking-systems-full-reserve-vs-free-banking/

(26) The Time Value of Lightning Network, Nik Bhatia & Joe Consorti. Available: https://theBitcoinlayer.substack.com/p/the-time-value-of-Lightning-network

(27) Money: At The Center Of Transactions, Ceyna Oder & Irena Asmundson. Available: https://www.imf.org/en/Publications/fandd/issues/Series/Back-to-Basics/Money

(28) The Time Value of Lightning Network, Nik Bhatia & Joe Consorti. Available: https://theBitcoinlayer.substack.com/p/the-time-value-of-Lightning-network

(29) Bitcoin Banking Systems: Full Reserve vs Free Banking, Eric Yakes. Available: https://yakes.io/Bitcoin-banking-systems-full-reserve-vs-free-banking/

(30) Circulation credit is that which emerges from fractional reserve banking under the Misesian framework. Explanation available: https://mises.org/wire/ludwig-von-misess-circulation-credit-theory-trade-cycle

(31) It is somewhat beyond the scope of this paper to go too far down this rabbit hole, but it ought to be mentioned that clearinghouses pose political problems in addition to economic ones: most concisely that the power they come to acquire over the financial system they lubricate make them ready attack vectors for the soundness of the money for which they are responsible. Modern central banks started as clearinghouses in some sort or another and universally moved to abandon even the pretence of "reserve" over a long enough period of time. Arguably, a precondition for this shift is in theory and has been in practice the greater commercial and transactional utility of bank fiduciary media than of specie. Hence the (potentially belabored) discussion above outlining that Bitcoin is significantly more transactionally and commercially useful in self-custodied contexts than gold is, or that any other form of base money has ever been. It could be argued that what makes the likes of fedimints, Bitcoin-backed eCash, and all manner of Bitcoin banking interesting in the first place is that it is not a commercial necessity the way fiduciary media was to gold, and which has led to modern-day fiat: it merely adds to a spectrum of potential trade-offs, yet is rooted in a bedrock of distributed self-custody that lacks the (now known to be) fatal fragility of gold.

(32) It is worth contemplating the likelihood that LSPs would likely establish private, commercial scale fedimints to act as de facto clearinghouses for their own liquidity rebalancing needs.

(33) The Web-of-Stakes scheme leverages information within the Bitcoin ledger to create transparency of one's economic stake and history of behavior within the financial system. The concept expands upon the idea of a Web-of-Trust, a decentralized trust model fundamental to modern cryptography. The Civ Kit whitepaper: https://github.com/civkit/paper#civ-kit-whitepaper

(34) A discussion of redemption attacks can be found in the writing preceding this: https://yakes.io/bitcoin-banking-systems-full-reserve-vs-free-banking/

(35) While he proposed a detailed scheme he was not the first to theorize the concept. The project Scrit (no longer active) was the first to propose the concept. Available: https://theblockchaintest.com/uploads/resources/NA%20-%20A%20distributed%20untracable%20electronic%20cash%20system%20-%202019%20-%20Nov%20-%20Paper.pdf

(36) A detailed overview of the proof-of-liabilities scheme can be found here: https://gist.github.com/callebtc/ed5228d1d8cbaade0104db5d1cf63939#file-ecash-pol-md\

(37) Calle's response to the author's initial review of this scheme available: https://gist.github.com/callebtc/7b60506343a7a3dc796e03144f0ed6f6

(38) However, the effect of which is likely to be minimal as worst case scenario (in the context of Cashu) the mint would become aware of the LN transaction sent to it to receive the exposed eCash and LN sender privacy is generally strong enough to not reveal the actual on-chain pseudonym of the sender.

(39) Image source is here: https://gist.github.com/callebtc/ed5228d1d8cbaade0104db5d1cf63939#file-ecash-pol-md

(40) It could also simply prove true that privacy costs do not matter to users and thus the adoption of this technology natively to wallets becomes close to ubiquitous.

(41) A more detailed description of Hoseki can be found on the company's website. Available: https://www.hoseki.app/

(42) A simplified overview of the transaction mechanics: https://twitter.com/_AlexLewin/status/1667185028768452611

(43) An overview of coinjoins: https://Bitcoinops.org/en/topics/coinjoin/

(44) An overview of channel factories: https://Bitcoinops.org/en/topics/channel-factories/

(45) The CFO of Visa has stated 65,000 TPS is the theoretical maximum of the network. Available: https://cointelegraph.com/news/bitcoin-lightning-network-vs-visa-and-mastercard-how-do-they-stack-up

(46) 1 bit is 1/1,000,000th of a bitcoin and 100 sats.

(47) For a more technically detailed discussion of the topic: https://delvingbitcoin.org/t/thoughts-on-scaling-and-consensus-changes-2023/32

(48) Available: https://lightning.network/lightning-network-paper.pdf

(49) Anthony Towns referred to forced expiration spam as the "thundering herd" problem. Available: https://lists.linuxfoundation.org/pipermail/lightning-dev/2023-September/004095.html