

Only The Strong Survive

A Philosophical, Technical, and Economic Critique of Prospects in “Crypto” Beyond Bitcoin

Allen Farrington & Anders Larson



photo by [Simone Pellegrini](#), via Unsplash, colorized by the authors.

v 0.2

Allen Farrington holds a board observer role at Blockstream. Anders Larson works for a US-based hedge fund.

*What follows **does not represent the views of either author's employer** and is **not financial or investment advice**. It is intended as a philosophical, technical and economic assessment of a novel class of internet protocols. These protocols mostly happen to give rise to natively digital assets, which lend themselves to naturally emerging online and effectively public markets, and which present direct investment opportunities. Nonetheless the following is merely **and only** our opinion of how these technologies are likely to progress. Readers considering investing in any asset discussed herein should **do their own research and should not rely on our work**.*

ABSTRACT

In this piece we discuss our concerns around the broader “crypto” space beyond Bitcoin. Our arguments revolve around three primary themes: Philosophical, Technical, and Economic. From a philosophical perspective, we discuss the core properties that make Bitcoin work and that make it unique, in our view. From a technical perspective, we evaluate how these properties are lacking to a greater or lesser extent in alternative “crypto assets,” having been designed out again on questionable philosophical justifications. From an economic perspective, we walk through our concerns that “crypto” has not shown a path to establishing a basis for justifiable real-world value, explain why we believe this would be very difficult to ever accomplish, and suggest therefore that their technical flaws are especially vulnerable. Given the complexity of the task we have set ourselves, this piece is intended to bring our arguments into long form and provide a discussion base for good faith disagreement.

“Talk, talk, talk: the utter and heart-breaking stupidity of words.”

William Faulkner on the authors of this piece

Bitcoin is an important innovation. We doubt any serious financial professional now disputes this. However, the importance of Bitcoin’s many imitations, collectively dubbed “cryptocurrencies” – or just “crypto” – is contentious. Our argument is that, in the long run i) it is likely that much of the current value will prove unsustainable and therefore will disappear and ii) it is possible similar systems built on Bitcoin will capture much of this value. This paper is a rigorous philosophical, technical, and economic analysis of why we hold this view.

Given we use the terminology throughout the paper, we must be absolutely clear on both what we mean by “decentralized finance” and “DeFi” and our attitude towards it. This is equally for the purposes of clarifying our chosen terminology and clarifying our motivation and attitude. We are strongly supportive of the principles of decentralized finance, which we will do our best to explain and examine shortly.

By *the concept* of decentralized finance, we mean an ecosystem in which the building blocks of financial and capital markets products are freely accessible to all without having to navigate technical bottlenecks or economic middlemen; in which their workings are transparently inspectable and auditable on the basis of free and open-source code; in which not only the products but even the architecture of marketplaces operates on these principles; in which this constitution lends itself naturally to programmability at origin and interoperability thereafter; and in which, due to the combination of all the aforementioned factors, no individual or entity can maliciously or politically affect market activity, be this in the form of agitating to advantage themselves or to disadvantage others. The dream is, in effect, that all participation is honest, and all honest participation is accepted, such that a long and inclusive tail can be unlocked for capital market activity – a domain of social and political economy notoriously inefficient, exclusive, extractive, and oppressive throughout modern history.

Our argument is that *this instantiation* of the concept is mostly built on unfirm foundations and, unfortunately, is unlikely to last. It is also worth being as upfront as possible that much (probably all) of the financial constructions we criticize could eventually end up being reconstructed on Bitcoin – on *Bitcoin DeFi* as we might call it – and would be equally worthy of criticism. Our argument here is not that Bitcoin and its higher layers are unequivocally good, but rather that such questionable constructions seem to be encouraged – even *necessary* – in the non-Bitcoin environment so as to fill a gap in justification of robust value that emerges in the first place *due to having injudiciously altered Bitcoin’s design in their own construction*.

Our opinion of the folly of these design choices makes up Section 1 of the paper. We therefore have two worries about non-Bitcoin crypto DeFi that builds on this folly, which make up Sections 2 and 3 of the paper, respectively: that the design folly means they are unlikely to last as allegedly robust, decentralized assets, and: that this fundamental flaw has implicitly encouraged a more superficial flaw of attempting to bootstrap an alternative foundation of value, but which, on rigorous inspection, cannot be justified (because the applications are not *really* finance). Section 4 evaluates the contemporary investment rationale, and Section 5 explores our belief that much, if not all, of the functionality will eventually be rebuilt on Bitcoin.

To clarify our terminology, therefore, by either “DeFi” or “decentralized finance” we mean only the concept. By “crypto” we mean the current, non-Bitcoin based instantiation. We appreciate this may cause some confusion given *DeFi* seems to mean both the concept and the instantiation in different situations, and *crypto* seems to mean public blockchains both inclusive of and exclusive of Bitcoin when spoken by different people or in different settings, but frankly, we had to bite the bullet somewhere. Given we want to emphasize (repeatedly!) our fondness of the *concept* of decentralized finance, this seemed like the best permutation on which to settle.

Overview and Section Summaries

1 - The Innovation From First Principles

We provide a philosophical analysis of what we believe to be the real innovation in Bitcoin: proof-of-work and the difficulty adjustment enabling distributed consensus and “money” as an emergent and endogenous use case.

2 - Crypto Is Not Decentralized

We argue that the variety of changes crypto projects make to Bitcoin’s design make it likely that all instantiations will have to centralize at some point, and in some form. This both undermines the alleged goals and makes it a questionably costly alternative to already centralized systems, but also introduces an even more ominous attack vector: not even technical or economic but instead social.

3 - Crypto Is Not Finance

We argue that the attempt to bootstrap value not by security but instead by utility – in particular, utility in financial applications – does not solve crypto’s core problem but only exacerbates it and delays its resolution. We also highlight that the common and popular metrics that capture the extent of apparent “financial activity” are deeply misleading as indicators of health and success and primarily serve to encourage further capital inflows that are, in fact, necessary to sustain the ecosystem, but without contributing to real economic productivity.

4 - The “Investment” Rationale

We provide a rationale for investing in the space to date, and argue that the most sensible investment thesis is essentially a subtle category error that results in transferring across principles from software venture investing that do not quite apply in this space. We argue

further that only a certain class of investors is likely to commit this error, and the realization that others will not follow will likely mark the beginning of the end.

5 - Layered Architecture and Gall's Law

We argue that the desirable features of DeFi will likely emerge before too long on Bitcoin. Furthermore, we argue that the fact of these features taking longer and being more difficult to build is a fundamentally good thing; it reflects that Bitcoin's architecture has been built in a more robust and prudent manner than its DeFi peers. Ironically, in the long run, this is likely what will enable extension of functionality. We give basic details on a handful of relevant projects before analysing this dichotomy in more philosophical detail.

6 - Why We Might Be Wrong

In this short final section, we give a (non-exhaustive) list of reasons why the above analysis may turn out to be mistaken.

Appendix A - Common Pool Resources

In Appendix A, we argue that all "crypto assets," Bitcoin included, are properly understood as "common pool resources," as opposed to, for example, public goods, or truly private property. We then argue that, according to arguably the most respected analysis of such entities, the governance characteristics of Bitcoin are excellent while those typical of crypto are mediocre to poor. This is not crucial to the main argument of the paper, but may be of interest to readers interested in academic political philosophy and economics.

Appendix B - Rehypothecation Algebra

In Appendix B we provide the workings for the discussion near the start of Section 2 as to how to *actually calculate* systemic exposure, contrary to naive claims of "overcollateralization ratios" common in the space.

1/n – The Innovation From First Principles

*One Ring to rule them all, One Ring to find them,
One Ring to bring them all, and in the darkness bind them.*

J.R.R. Tolkien on “Crypto”

TLDR: *in Section 1, we provide a philosophical and technical analysis of what we believe to be the real innovation in Bitcoin: proof-of-work and the difficulty adjustment enabling distributed consensus and “money” as an emergent and endogenous use case.*

Jump to [Section 2](#) for how crypto attempts to upgrade this innovation, but potentially just breaks it.

What is the innovation of Bitcoin? Is it a decentralized computer? Is it digital real estate? Is it a peer-to-peer payments tool? Is it a solution to [the Byzantine Generals Problem](#)?

Answering the first question depends on understanding the context in which it is being asked. The questions that follow, in response to the first, each suggest their own context. But we believe all to be unhelpful. We offer two contexts of our own, and hence two answers. Technically, the innovation is the proof-of-work algorithm, the difficulty adjustment, and the native monetary unit that is able to endogenously emerge. Socially, on the other hand, the innovation is an immutable and uncheatable distributed ledger.

Let us now break out the two categories: Technical and Social.

Technically: The proof-of-work algorithm and difficulty adjustment enabling distributed consensus on censorship-resistant, [integrity-assured](#), floating value *is the breakthrough*. What the design *achieves*, is a *social* matter, and in purely technical terms, nothing else about Bitcoin is either original or all that interesting. Public key cryptography and hash functions have been perfectly well understood for decades and implemented across industries and applications without much fanfare.

Achieving consensus is a compromise. Permissionless consensus involves an achievement of compromise previously impossible due to problems in coordination and incentives. Every actor may have an incentive to manipulate consensus to their own benefit. The innovation of Bitcoin is the creation of an incentive to protect honest contributions to consensus from a distributed network that is greater than the incentive to either attack honest contributions or to submit dishonest contributions. This way there is an economic incentive to reach consensus in a decentralized environment that otherwise lacks a single source of truth. This innovation is implemented via proof-of-work: value is programmatically escrowed and probabilistically returned on the basis of valid and honest contributions, and is otherwise confiscated. These contributions are both perfectly competitive and auditable. It is the work that is respected, not the worker. The worker need not be known to the rest of the network.

That the resulting protocol can be considered “money” is endogenous to this incentive scheme – but it is also necessary. The reward for honesty must have a denomination that can reliably be weighed up against the cost of escrow, dishonesty, or overt attack. This enables a native monetary unit. The difficulty adjustment enables cryptanalytic stability in the provision of security via work, the aggregation of work as input, transformed by the algorithm, will always create an output with the same temporal properties, enabling decentralized timing as well as decentralized truth. A verifiable order of events allows the updating of the ledger to be sensibly interpreted as transfers of balances of the monetary unit, and stability to inputs makes the incentive scheme flexible and reactive to real-world costs expensed by the growth, or shrinkage, of the network, enabling the value of the monetary unit to float.

This is the breakthrough that allowed Bitcoin to advance beyond previous attempts at digital cash, in which the hoped-for digital *value* was assumed to correspond to an external yardstick, almost always a given denomination of fiat money. Previous attempts disallowed a workable incentive scheme to protect the integrity of the asset in a distributed manner. Without such a scheme, value and timestamping had to be guaranteed centrally. This grounded that value, ultimately, in trust in the issuing authority. This made the asset, in whatever form it took, little more than an IOU in digital guise.

This breakthrough also allowed Bitcoin to advance beyond previous, effectively decentralized, schemes. Plenty of which exist. Arguably every internet protocol is one; from the web to email to TOR to torrents to git. Because they contributed to decentralized consensus and *immutability* via unforgeable costliness. Any attempt to manipulate the record by passing off a change as honest would require a single actor (re)performing as much *work* as the entire distributed network had to that point, and racing it in real time to overtake its aggregate. That this is practically infeasible grants Bitcoin the additional benefits of transparently verifiable accuracy of the timing of transactions. In addition to the high cost of attack via attempted dishonest contribution, contributes to a low cost of defending the network against such an attack.

The entire decentralized system needs the prospect of value to motivate the provision of censorship resistance and integrity assurance. Hence, the “value” realized by the network must be emergent and endogenous; it must *float* relative to any external measure.

Bitcoin’s distributed security is endogenous and depends on its value as money; and its value as money is endogenous and depends on its distributed security.

Socially: Bitcoin is not a computer, it is a ledger. This is not to deny Bitcoin is “money.” Of course, it has a monetary aspect, but only because money is arguably the most important ledger, in whatever form it happens to take. Adam Back essentially made this point [on Preston Pysh’s podcast recently](#), reflecting on previous failed attempts to create digital cash. One way of conceiving the real breakthrough of Bitcoin is that Satoshi realized stable value *could not be robustly targeted*. The robustness of the ledger itself would have to be paramount, from which value *might* endogenously emerge, if the prospects of the protocol design are appreciated by the market and speculated upon. This is as far as we need to go into the philosophy of “what is money?”, given that referring to money as a ledger is not terribly controversial.

What *may be* controversial is our claim that the presence of computers in the workings of Bitcoin [is a red herring as far as this analysis goes](#). Computers are useful because they allow the introduction of mathematical precision into the necessarily probabilistic *proving of work*. The presence of computers wrestles with the inherent uncertainty of distributed competition and reduces it to comprehensible and manageable statistics. This is useful for those weighing up monetary costs and benefits. Also, “the Internet” is useful because it allows the desired consensus to be reached and verified *globally* nearly instantaneously. But neither *technically* matter. Bitcoin could work in principle with pen and paper and carrier pigeons bringing the result of computations-by-hand back and forth to a public square: the *ledger* would work, but its intended use case as *money* likely would not. Computers make Bitcoin far more useful, but they do not *make Bitcoin*.

The consensus (i.e., the *ledger*) is what matters. The breakthrough of proof-of-work and the difficulty adjustment is far more about harnessing energy and existing communications infrastructure to calibrate incentives than it has to do with computation. Computation is a tool that makes the process more efficient and reliable, but Bitcoin is not itself a “computer.”

Calling Bitcoin an “inefficient database” is equally unhelpful. A car is an inefficient stove if you run the engine hot enough to fry an egg on the hood, but only because you are using it for something other than its intended purpose. The Bitcoin timechain exists as a digital data structure due to the coincidental presence of computers as just discussed, but not as an essential feature. Rather than a “database,” what matters is that it is an append-only ledger, appendable by anybody, and verifiable by anybody that it has only ever been appended. In other words, a ledger that is *censorship resistant* and *integrity assured*. This could be implemented by hand and then all discarded besides the UTXO set: an unordered list of mostly optical gibberish – hardly a “database.”

If you want to call a given implementation a database, knock yourself out, but don't expect that this will help you understand it any better. It would be like calling the HTML, CSS, and JavaScript on a web server a special variety of text document, or perhaps a form of atrocious blank verse. This is true, but tells you nothing, and is more likely to confuse you than it is to help you understand what is happening or why.

There are a handful of interrelated characteristics that constitute the real innovation of Bitcoin and that delicately balance to give it unprecedented functionality. These are:

- i) The proof-of-work algorithm
- ii) The difficulty adjustment
- iii) The native unit of (only) monetary value
- iv) The lack of a founder or acknowledged leader
- v) The economic incentive created for distrusting individual actors to achieve distributed consensus, unforgeably and immutably.

This all allows Bitcoin to realize endogenous value as an asset grounded in its security, and endogenous provision of security as incentivized by this asset.

Our thesis is that all non-Bitcoin crypto projects, usually in an attempt to add functionality deemed to be an improvement on that offered by Bitcoin or that is even fundamentally *impossible* to offer on Bitcoin, necessarily sacrifice at least one element just outlined. [Section 5](#) will argue that this thinking is impatient and misguided in the long-run and that the hoped-for functionality is likely to slowly but surely emerge on Bitcoin. Next, [Section 2](#), will explore what we believe to be the consequence of this sacrifice, which we believe to evidence short-termism: that decentralization is put at risk.

2/n – Crypto Is Not Decentralized

“Where I’m from, only the strong survive.”

Allen Iverson on “Crypto”

TLDR: *in Section 2, we argue that crypto’s design differences to Bitcoin, as discussed in [Section 1](#), make it difficult to live up to its promise of openness and trustlessness and hence likely if not inevitable that it will have to centralize at some point. This will likely both undermine the alleged goals and turns it into a questionably costly alternative to already centralized systems. It also introduces an even more ominous attack vector: not even technical or economic, but instead social.*

Jump to [Section 3](#) for a critique of the “financial” pretensions of crypto.

Any hope for *decentralized finance* must be grounded in a technology that is reliably and probabilistically secure, open, and distributed. We reiterate once again our belief in the potential of this *concept*, were it to be built on such a secure, open, and distributed base. Our problem with crypto – with *this instantiation* of decentralized finance – is essentially that it is not decentralized enough and it is not finance. In this section we will describe why we believe it is fair to say it is not sufficiently decentralized: why this is true in theory, how it manifests in practice, and how we believe it is likely to develop.

As far as we can tell, the intention behind the handful of ways this has happened is to attempt to improve upon some parameter of the Bitcoin timechain’s operation: its block time or regularity, its inflation schedule, its programmability, its privacy, the difficulty at the layer of the timechain to either introduce total token fungibility or definable nonfungibility (arguably a special case of programmability, but also, an enormously popular one) or more exotic goals as well.

To avoid the accusation that our criticisms are cherry-picked or that we are ignoring the benefits of what has so far been built, we will first do our best to steelman what is *gained* by making these changes.

In social terms, Marvin Ammori gives the following helpful characterization of the ecosystem in [Decentralized Finance: What It Is, Why It Matters](#):

“With DeFi, anyone in the world can lend, borrow, send, or trade blockchain-based assets using easily downloadable wallets without having to use a bank or broker. If they wish, they can explore even more advanced financial activities — leveraged trading, structured products, synthetic assets, insurance underwriting, market making — while always retaining complete control over their assets.

DeFi protocols abide by key criteria — in particular, permissionless-ness and transparency — reflecting values found in Ethereum, the open-source decentralized software platform that forms the infrastructure for most decentralized applications.”

He adds on the importance of permissionless-ness specifically that,

“The permissionless nature of Ethereum-based applications...collapses barriers to entry for entrepreneurs down to zero. End consumers are the primary beneficiaries of this innovative environment: Because all applications share the same database (the Ethereum blockchain), moving capital between platforms is trivial. This forces projects to ruthlessly compete on fees and user experience.”

This is an admirable achievement. It is easy to see why this combination of features enormously benefits consumers. Referring back to our initial characterization of the *concept* of decentralized finance in the introduction, *no individual or entity can maliciously or politically affect market activity, be this in the form of agitating to advantage themselves or to disadvantage others.* If Ammori’s account is accurate (now and indefinitely into the future) then clearly any malicious advantaging of oneself, or abuse of others, will be met with immediate and unstoppable abandonment of the adversarial environment. It stands to reason that this possibility should keep participants honest in the first place.

Balaji Srinivasan gives an interesting technical accompaniment to the social interpretation just presented, arguing in [Yes, You May Need A Blockchain](#), that, *“public blockchains are massively multicient databases, where every user is a root user.”* By “root user,” Srinivasan essentially means, in straightforward English, somebody with the rights and permissions to change whatever they want about the structure of the database, rather than “regular users” (clients) who can only read from and write to the database within prescribed rules.

Srinivasan’s point in establishing this somewhat technical premise is to make the following argument, which is rather powerful,

“Different applications typically don’t ... give users certainty that their data wasn’t intentionally tampered with or inadvertently corrupted during all the exporting and importing.

The reason why boils down to incentives. For most major internet services, there is simply no financial incentive to enable users to export their data...some call this the data portability problem, let’s call it the data export/import problem to focus attention on the specific mechanisms for export and import.”

This underpins Ammori’s point above, that users can *“retain complete control over their assets”* and *“[trivially] move capital between platforms.”* This can happen because all applications point to the same underlying “database” in which all users are root users: no user has the privileged ability to override the activity of others.

What allows this in the first place, Srinivasan argues, comes down to *incentives*. Regular databases have no incentive to enable seamless import and export, and possibly their operators have social or corporate *disincentives* in addition. “Public blockchains” on the other hand, have a fundamentally different feature that allows these incentives to exist. Srinivasan explains,

“Because the data associated with a public blockchain represents something of monetary value, it finally delivers the financial incentive for interoperability. After all, any web or mobile app that wants to receive (say) BTC must honor the Bitcoin blockchain’s conventions. Indeed, the application developers would have no choice due to the fact that Bitcoin by design has a single, canonical longest proof-of-work chain with cryptographic validation of every block in that chain.

So, that’s the financial incentive for import.

As for the incentive for export, when it comes to money in particular, users demand the ability to export with complete fidelity, and very quickly. It’s not their old cat pics, which they might be ok with losing track of due to inconvenience or technical issues. It’s their money...Any application that holds it must make it available for export when they want to withdraw it, whether that means supporting send functionality, offering private key backups, or both. If not, the application is unlikely to ever receive deposits in the first place.

So, that’s the financial incentive for export.”

Srinivasan builds toward the conclusion that,

“This is a real breakthrough. We’ve now got a reliable way to incentivize the use of shared state...while enforcing a common standard and maintaining high confidence in the integrity of the data.

This is very different from the status quo. You typically don’t share the root password to your database on the internet, because a database that allows anyone to read/write to it usually gets corrupted. Public blockchains solve this problem with cryptography rather than permissions...

In other words, public blockchains are massively multiclient open state databases where every user is a root user.”

A fundamental question of all that has been described above: Does any of it require a token? Recall tokens are by no means the essence of “decentralization” – HTTP, email, BitTorrent, TOR, Git, and wikis are decentralized and involve only exchanging information, not value.

Srinivasan argues it does because he recognizes that the technical blueprint borrowed from Bitcoin necessitates a native unit of value. But any self-respecting representation of value, or system purporting to guarantee representations of value, must possess precisely the properties teased out in [Section 1](#): there must be consensus as to who owns what, for starters. If this is to be achieved solely in the digital realm, where only information exists and where the transformation of information is effectively free, we need a means of making dishonest claims to ownership or transfer of ownership both identifiable *as dishonest*, and, preferably, disincentivized on the basis of being too costly to attempt in the first place. We require immutability and unforgeable costliness.

Which is all to say we presuppose a distributed ledger with censorship resistance and integrity assurance. But imagine then that the token is intended to be something other than money. We anticipate three conceptual issues with whatever design is then conceived, all of which will threaten any previously credible claims to “decentralization”: i) it will lose a fight for liquidity with *actual money*, ii) it becomes a poor economic signal for coordinating security provision, and iii) the timechain itself bloats.

Everything Fights for Liquidity: why would a non-monetary token need to *have a nonzero holding period* and hence a non-negligible value in the first place? In short, why would anybody hold something that is *money, but only for X*, for any period other than the nanosecond required to transact it? It’s analogous to the choice between dollar bills and casino chips of the same denomination. Why willingly take on the casino’s liabilities *and* restrict the ability to engage in any other commerce? Why not buy casino chips when and only when you go to the casino?

This isn’t a death knell in and of itself but note it relies on a far more delicate consensus than would an asset straightforwardly trying to be money, and nothing more. If you can’t rely on others holding it (because it’s money and nobody needs a particular reason to hold money) then sustaining its value very likely requires some kind of coordination. And worse still, whatever coordination is arrived at likely has no inherent economic incentives that independently motivate coordination (as in, besides the attempt to realize benefit that *depends on* the coordination) *and* must overcome the economic disincentive to just hold money instead. In other words, everything fights for liquidity.

As far as we can tell, the current “independent motivation” to fight for liquidity and hold *most* crypto assets for a nonzero period is “yield.” This is a good segue to [Section 3](#) given we have arrived at requiring that the “use case” is, in fact, that the tokens are to be considered securities:¹ their value is realized via the rights to future cash flows. But we will leave this for now and merely comment on the inherent contradictions already engendered given i) a security, and a yield, assumes and requires a separate money, and we set out to architect

¹ We apologize for the potential for semantic confusion here but it is regrettably unavoidable. By “security” we sometimes mean “robustness of the timechain and ability to defend against attack via economic incentives in a decentralized and uncoordinated manner, and sometimes, as here, “asset entitling owner to the stream of future cash flows produced by some defined enterprise”

all of this so as to both only require itself and yet not be money, and ii) a “yield” is not a “utility.” This is a dire philosophical error: utility can be immediately realized, yield cannot; utility is non-monetary and is priced, yield is defined as a ratio of a flow of money over a stock. But we will come back to this – for now, we will leave our argument as: if one tries to win the fight for liquidity by toying with the very concept, we doubt one will achieve anything productive.²

Poor Economic Signal For Coordinating Security Costs: Energy has a real cost. Proof-of-work and the difficulty adjustment provide a way to manipulate the *effectively free* activity of transforming information – the only material in the digital realm – into an actually measurably costly process, hence providing a link between the physical and the digital. If you can’t rely on endogenous value as money, there will be no sensible basis on which to evaluate the merits of the real cost of security contribution. We don’t want to overstate the importance of this point: obviously, it *can be done* if the tokens have value for any reason whatsoever. Our suspicion is rather that the process of justifying this decision is far less of a rational calculation than were the security reward nothing more than endogenous monetary value; it bakes into the calculation an element of belief (we might say a disregard for the robustness of this link between digital and physical reality) that there is no particular reason to expect to be stable or reliable over the long-term. Once again, we arrive at a coordination problem, the resolution of which does not benefit from any inherent economic incentive; the solution to the coordination problem is precariously rooted in the desire for there to be a solution.

Money *is just a type of information*. More specifically, it is the information reflecting the social consensus of the value of *time*. Bitcoin *is just a ledger* of transfers of value. As a form of money endogenous to distributed consensus, the information in the Bitcoin timechain is arguably as pure a representation as possible of precisely what costs and what time have been contributed to its security. This leads to Bitcoin’s *continued existence and honest and valid distributed operation*. Were the token intended for some economic property besides pure information, or too far removed from information, then it is likely the signalling mechanism linking the information to the cost of its provision and security and the value of time will at some point start to falter.

Timechain Bloat: this is by no means *logically* necessary, but it is easy to predict as a possibility and easy to observe as it has happened often. If a timechain is structured so as to contain more than this bare minimum of information (either in the form of economically loaded content that takes up far more data than validation of monetary balance transfers, or just *too much validation*) then it may reach such a size that it becomes practically or economically impossible for many to either run a node or contribute to security. This could be either in absolute terms or in the throughput of keeping a live view of timechain validly updating. It is impossible to reason conclusively about such an outcome entirely in general, as any touted thresholds will be arbitrary. But there will be some point at which this is concerning, given how degraded immutability and unforgeable costliness have become. Whether the community input on any technicality of the consensus mechanism remains sufficiently decentralized also risks becoming jeopardized.

It is easy enough to predict that, as these interrelated problems start to make themselves felt, so too will the pressure to ward off security issues by further centralizing control of the timechain. We repeat this does not follow necessarily from our argument thus far – it is in no sense predetermined – it is merely unfortunately easy to articulate as an increasingly likely possibility. Ammori draws attention to this,

“The underlying backend infrastructure for DeFi, Ethereum, must continue to scale in order to support higher bandwidth demands. Processing approximately 1.5 million unique transactions per day, Ethereum is already at its current max capacity, and transaction fees have spiked as a result.”

But notice “transaction fees spiking” is *good for security!* So we have a somewhat perverse situation in which the more secure the protocol becomes, the more its value proposition suffers. In order to become “more usable” it has to become less secure.

² A variation of this point was originally articulated in John Pfeffer’s, [An \(Institutional\) Investor’s Take on Cryptoassets](#), which, to our minds, has never been adequately rebutted, and there is a lack of acknowledged resolution to the existential concerns raised.

This is all arguably exacerbated by proof-of-stake, a consensus algorithm proposed in contrast to proof-of-work in which economic incentive is provided in the form of locking up capital for the right to validate blocks, having it stripped if deemed by the community to have done so dishonestly, but otherwise being rewarded with some combination of newly issued tokens and more manageable fees. But notice the tacit admission of having lost the fight of liquidity: the basis of value grounding has *explicitly* become “yield,” meaning the token is fighting for liquidity not on the basis of being money, but of being a security. Of course, it is not marketed as such, but we believe the economic logic here is clear and becomes remarkably easy to discern once recognized. In a blog post titled [Why Proof Of Stake](#), Ethereum co-founder Vitalik Buterin even refers to the validator reward under proof-of-stake as “interest.”

To be clear, money does not bear interest: securities do.

The core philosophical issue at play rests on the answer to what was left dangling as a rhetorical question: *do you need a token for that?*, which seems to be, *no, you do not*. Because “that,” whatever it is, is not money, hence it will lose the fight for liquidity, it won’t sustain endogenous value, nor will it sustain endogenous security. For the tokens of a timechain, if value is exogenous, then security must be exogenous. If security is exogenous, then value must be exogenous. Without pure and simple “money” – concise information on the social value of time – both value and security will likely have to end up exogenous and coordinated to a greater or lesser extent, at which point we have sacrificed decentralization almost by definition.

It seems to us that, from a purely philosophical perspective, the only way to stop the technical and economic structure of a non-Bitcoin timechain from breaking is therefore to effectively centralize it in some or other way. This brings us to a curious junction because i) whatever we are talking about would clearly no longer be *decentralized* finance, however impressive it otherwise may be, and ii) this centralization is likely to introduce an entirely new attack vector of *whoever is in charge*. Surely this makes the ecosystem not only *not decentralized*, but in fact, incapable of living up to the promises of an actually decentralized system in theory as well.

Our rebuttal to the Devil’s Advocate of what is *gained* by non-Bitcoin crypto is therefore quite simply that what is *lost* is the credible ability to resist attack. The fanciest cryptography facilitating the most complex securitisation scheme with the fastest confirmation time the world has ever seen will be worth precisely zero if or when the infrastructure supporting it is attacked in such a way that it can no longer claim to be secure, open, and distributed, at best, or simply ceases to exist, at worst. To be clear on this point, price activity in the meantime is more or less completely irrelevant. If anything, positive price activity attracts attack, be it from economic speculators, technical exploiters, or state based social actors.

Why this obsessive focus on security, vulnerability, and costs to attack and defend? This might be a strange way to talk about what Ammori calls a, “*globally accessible supercomputer*,” whose, “*native programming language (Solidity) can be used to create any conceivable application*.”

Because timechains are not fundamentally computers. They are (or ought to be) ledgers. We see crypto as facing an impossible dilemma: if the timechains are computers in any sense, they are surely only valuable as *decentralized computers*. Ethereum is *much* more expensive than AWS, [by just about every conception of cost](#): write cost, storage cost, write speed, and sync speed. If, at the end of all of this, crypto protocols turn out to be every bit as centralized as cloud computing, the debate ends precisely at the point on which they are outcompeted on cost by many orders of magnitude.

The benefits *must be* some feature that follows from *decentralization*, which in turn *must be justified* on the basis of superiority to AWS. What benefits are afforded the user of Ethereum that cannot be derived from AWS? Presumably some measure of censorship-resistance and integrity-assurance in their decentralized computation. After all, AWS can turn off its customers’ access, so the higher cost must surely be to overcome this danger.

Surely, following this logic, the user has to *know* that this danger has been overcome. Why pay the cost of uncensorable computation for censorable computation? This is why we focus on security, vulnerability, and the costs to attack and defend – because it is required for the decision to be remotely economically rational in the first place.

The vulnerabilities are far more than merely technical or economic. That sufficient technical or economic vulnerabilities exist at all create additional social vulnerabilities. “Cost” can be expended via bribes or threats to physical wellbeing. How much would it cost to bribe the possibly three or four Amazon employees who could literally flick the switch on Infura, [given ~22% of fully synced Ethereum “full nodes” are hosted on AWS, and ~71% of all fully synced nodes are hosted with some or other cloud provider](#)?³ And should any crypto protocols become genuine *financial infrastructure*, hence capable of being shorted if self-respecting in their ambitions, how good an investment would such a bribe turn out to be? Or, as Anthony Pompliano [recently put it](#) while interviewing Jack Mallers,

“All this stuff that claims to be “decentralized,” I just ask the founders, if the government came and said you had to go to jail if you don’t shut it off, would you shut it off? Oh, you could? Then it’s not decentralized.”

It is effectively free – or perhaps the cost is socialized – for the SEC or FATF or whoever to send a cease-and-desist letter. Secure, open, and distributed finance must be able to resist all of these attack vectors or it is pointless.

We do not mean to be morbid for the sake of it. We mean to be morbid because if or when *all the securitized value in the world* is at stake, there will be sophisticated attacks. These things *will happen*. If the intention is to *rebuild finance in a decentralized manner*, these attacks need to be rigorously studied and forestalled.

There is the perfectly realistic possibility that much of contemporary crypto becomes fully institutionalized. Not simply the involvement of institutions, which would itself be a necessary eventuality to a bull case, but something more like, *the total capture by institutions*. The probably intractable cost/attack dilemma could be solved by centralized and authorized hosting, which may also lead to shifting all operations to dollar stablecoins as a default. The outcome of something along these lines would transparently *not* be decentralized, open, nor any similar stylization that is touted as the advantage in this space in the first place. This outcome will not be mentioned again. For the purposes of our discussion, “institutional involvement” in crypto will mean that crypto is still *attempting* to be decentralized, and *attempting* to compete with legacy finance.

And yet, crypto has not been institutionalized in this way *and* seems to be attracting increasing institutional flows. It is eating Wall Street rather than being eaten and the dire consequences of centralization suggested as likely in the long run do not seem to have been fully borne out. Are these scenarios too pessimistic?

We think not, for two related reasons that form the basis of Sections 3 and 4: in [Section 3](#) we posit that what is actually *happening* in crypto reflects at least an implicit awareness of some of the problems outlined in Sections 1 and 2, and attempts to overcome them by bootstrapping its own grounded value in the form of financial utility rather than as money. Put another way, *as the virtuous circle of security and value from which money can endogenously emerge*. We argue this is doomed to fail in the long run as the attempt at bootstrapping does nothing to address the philosophical and technical issues that have been created by all attempts so far at moving the design away from Bitcoin. This is an important distinction: the argument is not that it is theoretically impossible to create a superior design of a public blockchain to Bitcoin. This is strongly suspected to be true, but it can hardly be *proven*. What can be observed and understood beyond any doubt is that no real-life attempt has come close to succeeding. Bootstrapping may actually introduce *more* such issues in their stead, this time

³ Due to timechain bloat and poor security signals.

primarily *economic*, and in either case makes the ecosystem only more systemically fragile. It also perversely delays the stress sufficient for revelatory fracture.

In [Section 4](#) we address that these early institutional flows are a kind of self-fulfilling prophecy with no happy ending: the flows prevent the resolution of the fundamental flaws, and provide more dry powder to create financial constructions that provide paper returns but solve no real problems. This attracts more flows, and perpetuates the irresolution of these flaws, maintaining and growing the illusion of secure, open, distributed, *and yielding* financial applications that are, in fact, nothing of the sort.

3/n – Crypto Is Not Finance

“There are three ways to make a living in this business: be first, be smarter, or cheat.”

Jeremy Irons as John Tuld in *Margin Call*, on “Crypto”

TLDR: in Section 3, we observe that non-Bitcoin crypto seems to have positioned itself as bootstrapping its own value by providing financial utility rather than money, and argue that this is likely fundamentally unsustainable and does not solve the core problem. Furthermore, we argue that most conceptions of the health of the ecosystem can only be sensibly interpreted as transient – be they “valuation” metrics, liquidity and solvency assurances, or the prospect of a link to real economic productivity.

Jump to [Section 4](#) for a critique of the typical rationale for investing in the space, as far as we believe we can identify it.

We reiterate our support for the concept of *decentralized finance* in general, and that our objection to contemporary crypto is not its purported aims, but the method of its attempt to achieve these aims. If anything, we feel the instantiation is letting down the concept, even if not widely realized at the moment. In [Section 2](#) we explained our view that the technical design choices inadvertently threaten any element of decentralization that is surely necessary for passable “decentralized” finance. In this section, we evaluate the economic merits (or demerits) of the “finance” this system realizes.

We believe that most crypto projects have attempted to escape the issues explained in [Section 2](#) by catalyzing value via utility instead of security, so far mainly opting for utility in *financial* applications, and hence adopting the moniker “DeFi.”⁴ We believe this is likely to end in tears because the base layer is neither money nor secure. The “finance” being constructed seems to us to therefore have the perverse effect of compounding an original fragility that follows from technical and economic unsoundness, *but also* buying time in the minds of those yet to become fully cognizant of all this. A charitable characterization might be that the hope is to bootstrap value via utility; a harsher characterization would be that it is all an attempt to borrow its way out of debt: monetary *and* technical. In short, whatever it is, crypto is not finance.

To lay some conceptual and rhetorical groundwork for what might seem like a sweeping claim, consider the expression “yield farming” – a concept that has driven enormous interest and capital into the crypto ecosystem. The concept does not refer to any real “yield.” A *yield* is the generated *flow* above maintenance or depreciation of the carrying capacity of some *stock* of economically productive assets. Less the recouped seeds for the next year’s crop, a harvest is a *yield* from a sown field. Less the financing costs, the interest on a bond is a yield. If the issuing business is solvent and profitable in unit economics-terms and hence the par value of the principal is relatively assured, the market will settle on a value that implies a probability of all the interest being paid as promised. The market assesses the productive carrying capacity of economic *stock* generating the ability to pay the *flow* of the interest.

So what *yield* is being *farmed* in crypto? There is transparently none. There are flows, but they are not *generated* by economically productive assets over time but rather appear near instantaneously as a result of speculative pricing across non-productive assets. The word “speculative” is not a denigration. There is nothing wrong with speculative value. But there is something bizarre and circular about discrepancies between different speculations on the potential future value *itself* forming the basis of profitable arbitrage that is then mislabeled as a “yield.”

⁴ A reminder to the reader that we will continue to say “crypto” for this instantiation and “DeFi” only for the concept precisely to avoid the obvious confusion this invites.

Perhaps this is just semantics? Redefine it not as a “yield,” but arbitrage typical of any market-making activity. Normal market-making itself relies on trading between those who naturally disagree in their assessment of speculative value. But the speculative value *of what?*

Of yields! This is the essence of *speculation* as opposed to, say, *appraisal*. Most financial assets are assigned value on the basis of their discounted future cash flows – that is, at least psychologically and philosophically if not mathematically. Hardly any financial assets of note have no capacity to generate a yield and simply have a value in the moment that a potential buyer believes to be mispriced – at least hardly any with noteworthy capitalization that makes them relevant to broader capital markets activity. “Speculation” arises from the inability to *know the future*. Yields take time and skill to generate. No real yield is generated instantaneously by arbitrage, and hence no real financial asset attains value and pricing this way, either. At root, financial assets must derive their value from a spot appraisal or from the prospect of future productivity from a stock of carrying capacity generating a yield. Does crypto?

No. Curiously, this is the simple part of the answer. Understanding what is happening is often a lot more complicated. The remainder of this section breaks into the following subsections to trace what is quite a complicated line of reasoning.:

i) an analysis of how the crypto ecosystem embraces and encourages rehypothecation, leverage, and securitization, liberally borrowed from traditional finance but that don’t quite serve any coherent purpose in this environment.

ii) how these three collectively contribute to creating systemic fragility such that, in general, capital has to keep flowing in to keep the system seeming healthy, and,

iii) an analysis of how the popular and common metrics worryingly both conceal the fragility analyzed in subsection i) and encourage the continuation of systemically necessary capital inflows as analyzed in subsections ii) and iii).

i) rehypothecation, leverage, and securitization

The first thing we want to highlight is the immense amount of rehypothecation of assets happening in the crypto ecosystem today. What we mean by “rehypothecation” is quite simple, albeit a little different to traditional finance (hence creating the potential for a great deal of confusion we also hope to dispel): a given asset can be used as “collateral” in one protocol, contributing to a new asset being *minted*, and then either itself reused, its collateralized “end product” reused, or its securitized governance rights reused, all again and again throughout multiple different protocols. It is actually difficult to disentangle these three separate concepts, hence us treating them all at once. Leverage on its own might not be so bad were the “rights” to the levered “end product” not dubiously securitized. Nor would securitization on its own be cause for concern were the securitization tokens not used as “collateral” and unboundedly rehypothecated.

Let us consider an example of how assets in the crypto ecosystem can be, and are, reused and recycled:

1. To start, a user takes \$1500 of Ether, deposits it into Maker, and gets \$1000 of the DAI stablecoin in return. This assumes a 150% collateralization ratio, a dubious metric, but it will do for now.
2. The user then deposits the new DAI as well as \$1000 of Tether into the Curve 3 pool and becomes an LP with \$2000 total staked into the pool’s liquidity
3. The user, as an LP, on top of garnering fees, is also granted a CRV token for being an LP. The CRV token was issued by the protocol as a “governance” token. The value of the CRV is supposed to derive from voting rights, over such matters as enabling access to treasuries or any other fees generated by the protocol. As of today, and in most crypto protocols, the “treasuries” of these protocols are actually just these self-created tokens.

4. This CRV token can then be lent out by the user using a lending protocol such as Aave. In Aave, the user can deposit these CRV tokens they received and earn interest on them, or collateralize against them for stablecoins in yet another protocol.
5. Once the “stablecoin” is generated from the CRV the user just used as collateral, the user is free to reuse these “stablecoins” to make another round of “investments,” in theory now going back to step 2 above.

Let us analyze Curve and 3pool in a little more detail. Curve is an automated market maker that facilitates the transfer of different Ethereum-based assets between users looking to trade. 3pool is a pool for stablecoins which allows users to exchange DAI, USDC and USDt and rewards participants with CRV, the Curve “token.”

Liquidity providers in 3pool – users with balances of stablecoins – deposit proportional numbers of DAI, USDC and USDt to provide liquidity for these exchanges. In return for staking these assets, the liquidity providers are paid back in two different forms. There is a “swap fee,” which can be thought of as users or “takers” paying a fee to the liquidity providers: stablecoins in this instance. There is also a staking reward to which liquidity providers are entitled. In exchange for locking up the underlying assets to the Curve protocol, the liquidity providers earn a certain amount of CRV in reward. The size of the transaction fee earned also scales with the amount of value staked. While clearly not advertised as such, it is fair to say that these staking rewards are effectively paid via unbacked seigniorage, which the protocol uses to incentivize makers to provide liquidity.

Next to touch on the native token to Curve. Why this token ought to have a value, and how it ought to be used, is unclear. The token has a programmatically unlimited supply and a vague value proposition for any investor. The governance to / claim on any “real” cash flows by any definition is unclear, even if defined so loosely as to potentially consist of receipt of other Ethereum-based tokens. On top of this, the only reason anybody would continue to hold their CRV tokens, from a capital allocation perspective, is a mechanism that gives more “rewards” in trading fees for continued staking.

Does that not sound rather convoluted? We are talking about a token whose only evidenced utility and function is to be re-staked into the very system that creates it *to create more of it*. The red herring of “governance rights” seems only to entitle the holder to contribute to the “governance” of this process of creating more of itself, and nothing more.

This clearly depends on new capital coming into the system to be sustained. But what happens when, one day, people decide they would like to own an alternative asset that actually provides access to real cash flows, or that has some more persistent basis for its spot valuation? Our contention is a “bank run” cascades the value of CRV down, and with it will crash the value proposition of running liquidity pools. Crypto investors who currently own CRV do seem to largely understand that, today, much of the governance is unproductive but hope that in the future the protocol will continue to mature and these tokens will accrue value as derived from their governance rights. Clearly, we are skeptical.




Figure 1: Explanation of why CRV has value

We use Curve as an example because it simultaneously encapsulates multiple risks we discuss within one subsystem of crypto. Curve enables the trading of Ethereum-based assets for each other, many (possibly all) of which have a market value in the first place that is justified via exactly the process we are currently analysing, in which the underlying assets are referenced either to each other's values, or to governance rights over flows of these assets (presumed to have value, hence the rights are presumed to have value). And as the icing on the cake, Curve created its own such token, the only utility or function of which is to re-stake into the same system, thus providing yet more pseudo-leverage to the Curve token ecosystem, and to crypto as a whole.

One further layer to all this is that many other crypto protocols aggregate user funds to help them become liquidity providers in Curve. An example is a protocol called Yearn Finance, which aggregates user funds and deposits it into protocols such as Curve to do the work of automated market making and entitle the original providers of funds to interest (styled as "yield" but we know better). Suddenly we have a crypto protocol (yEarn), that is built to allocate funds to other crypto protocols (Curve in our example), with users earning interest from the fees and once again the Curve token, CRV, but now also YFI token. From a user perspective just looking to garner "yield," YFI is a brilliant offering. How the yield is generated and what it actually represents, as well as the underlying risks is what concerns us. Funnily enough, there are even other projects that then attach onto Yearn Finance, notably Alchemix, but at this point we think our concerns have been voiced. Let us just say that we see no practical reason this iterative process would ever *need* to end ...

In all, what we have seen is how starting with \$1500 of ETH and \$1000 of Tether can lead to multiple different steps of assets being created and recycled, either via an implicit form of leverage or an implicit form of securitization. We (re)use the word "implicit" as these activities are not openly described as either "leverage" or "securitization" within the ecosystem, yet they seem to be the closest accurate descriptions in terms of traditional financial concepts: some of the protocols involved (MakerDAO above) take "deposits," of sorts, and mint new assets (more suggestively, *credit instruments*) determined to have some value as a function of the value of the deposit. This is essentially the role of a bank in creating *leverage*. Some other of the protocols involved (Curve above) mint new assets (more suggestively, *securities*) on the basis of the value ascribed to control of some or other resource. This is essentially the role of investment banks (or capital markets at large, but this is possibly just semantics) in *securitization*.

It is worth emphasizing that rehypothecation, leverage, or securitization are not necessarily bad things if employed properly and transparently. Our concern is effectively that they are being employed improperly and opaquely, such that systemic risks are created and which seemingly go entirely unnoticed, or perversely, the effects of which are misunderstood and celebrated!

For example, crypto proponents will often cite "overcollateralization" as a reason to be reassured that things can only go *so wrong*, or that, if things go wrong, we can be relatively sure the various structures in place can be unwound safely and the originators of capital made whole. To be clear on terminology before we get into the weeds, by $x\%$ *collateralization*, we mean every \$100 of synthetic asset is backed by \$ x (i.e. $x\%$ of \$100) of collateral. By $y\%$ *overcollateralized*, we mean that a synthetic asset is $100+y\%$ *collateralized*. We will try to stick to the former to avoid confusion unless it cannot be helped.

The idea that overcollateralization grants safety might be nice were it not for a naive arithmetic glitch in the reasoning just presented: only 200% collateralization (or greater) can achieve this *systemically*. For any lower ratio, *there will be* some number of iterations of rehypothecating collateral such that the value outstanding ex-initial collateral is greater than the initial collateral. Let's use the simple example of 150% to demonstrate:

\$150 (of actual dollars) is required as collateral to issue \$100 of StableAliceCoin, and likewise to issue \$100 of StableBobCoin. If we take our \$100 of SAC and pledge it as 150% overcollateralization to issue \$67 of SBC, our gross synthetic exposure is already \$167 on only \$150 of "real" collateral. In fact, we can keep on iterating this, pledging SBC to issue more SAC, and back and forth indefinitely, and approach a limit of \$300 aggregate synthetic value. 175% collateralization would take 3 iterations and only approaches a limit of \$233. 200% will never be overtaken by synthetic value in excess of the collateral, and the limit of unboundedly many

rehypothecations approaches \$200. ([Appendix B](#) elucidates the algebra from which all numbers presented in this discussion pop out).

What all this demonstrates, beyond any individual example, is that the idea of “overcollateralization” means something very different in isolation on the one hand, and in an environment in which assets can be endlessly rehypothecated on the other (“money legos,” one might say). What we arrive at is really *pseudo-leverage*. Nobody thinks of themselves as having put capital at meaningful risk, *because everything is overcollateralized*. And yet the ecosystem as a whole *is undercapitalized*. Far from enabling greater transparency, security, and so on, this connectedness of potentially individually robust elements perversely creates systemic vulnerability.

This may seem oddly familiar to fractional reserve banking (“FRB”) with the algebra causing déjà vu, to boot. What we have termed the percentage of “overcollateralization” could equally be thought of as the “reserve ratio” should the reader wish to translate all the algebra for themselves. But of course there are numerous drastic differences, the very fact the comparison can be made and that the algebra translates ought to set off alarm bells. We see three obvious issues here, although of course there may be more.

First, fractional reserve deposits are not “collateral,” but rather are savers’ funds lent at risk and put to productive ends. The “end” in crypto is not as of today productive but is rather a continued, and seemingly endless, recombination of the instruments. The purpose of *real* collateral is precisely that it is *not* rehypothecated but is used to guarantee the value of a securitized debt contract, given the obvious possibility that the loan itself will fail *on account of some real-world economic enterprise failing*. In crypto, none of these financial primitives are present, making it entirely unclear what the purpose of “collateral” is in the first place - or perhaps, more harshly, if the concept of “collateral” even applies.

Secondly, the purpose of the reserve ratio in FRB is to provide the liquidity of an equity buffer to manage the risk of maturity mismatch. It is not an entirely arbitrary leftover of recombinatory rehypothecation. Again, “maturity mismatch” is, in the first place, a product of the real economic activity FRB is intended to facilitate. Liquid reserves have a maturity of effectively zero and a definite value, whereas illiquid economic projects requiring debt finance have longer-dated maturity and indeterminate value, pending real world economic uncertainty. Given the lack of these primitives, it is unclear what the purpose is of the similar construction.

Lastly, in FRB, there is a lender of last resort in the form of a central bank. Not only is this a bad thing and not worth copying, it is what Bitcoin fixes. It isn’t even really possible to mimic in this environment in the first place. As a result, the role is taken up by pseudo-equity issuance to boost the equity buffer that also doesn’t quite make sense here. This is accomplished in the form of additional - essentially centralized - securitization, and additional primary capital from investors.

ii) systemic fragility

Given how much of the ecosystem is “collateralized” by Ether and other similar assets, some of which themselves generate new tokens, and given there is no clear link to stocks of economically productive assets (despite the prevalent securitization, underlying cash flows is not what has been securitized, to date) there is a need for real, external capital to act as the backstop of value. Taking into account the expected “return” this capital seeks, there is arguably also a need for continual *new* capital. Crypto has seen major drawdowns, but the bleeding has always been stopped by fresh capital being injected into the system. We suspect this has been largely due to the cheap capital environment across the world, and, perfectly ironically, the depression of real *yields* across alternative allocations, but we will cover this in more detail in [Section 4](#). After all, if you buy into the investment hypothesis anyway, then a 40%+ drawdown in 24 hours, as Ether for example experienced in March 2020, surely seems like an excellent opportunity to buy on price weakness.

Below we will speculate as to what might follow a similarly dramatic drawdown in the price of Ether, today or in the short to medium term future, and what this might mean for the health and wealth of the ecosystem at large. We should stress upfront that what follows is in no way scientific, nor is it a prediction we insist is predetermined. We readily admit the obvious criticism *that this is merely speculation*. This is true. Indeed, we contemplated including a historical analysis instead but found (obviously, on reflection) that no period and no data suffice given what we are speculating about here has never happened, to date. What interests us is that the thought experiment provides a means of better appreciating the extent of what we believe to be the systemic fragility caused by the not-widely-appreciated combinations of rehypothecation, leverage, and securitization described in the subsection above.

What would likely follow a large enough crash in the price of Ether is that protocols like Curve and others require more “collateral” to be put up or for loans to automatically close. Once this process starts, Curve tokens, which could be collateralized elsewhere, would likely begin to lose value. As this “bleeding” would play out, it is natural (and matches with the young history of the space) to expect some forced selling by market participants who otherwise believed their collateralized positions were relatively secure. With this forced selling of CRV and other tokens we would expect to see “Total Value Locked” and many other metrics used to describe the “value” of the networks begin to drop precipitously. This in turn would undermine the perceived value of the networks as well as the actual value of assets previously contributed, leading to anything for which their securitized “governance” tokens were used as collateral to also be subject to various liquidations and collateral calls.

As these liquidations and collateral calls ripple through the ecosystem, there would only be two ways the bleeding could end. Either nearly all of the leverage in the system would need to be wiped out, which of course would be a catastrophic decline in aggregate value, particularly so given our outline above of *just how much pseudo-leverage* can exist globally without anybody being locally aware. The alternative option is simply that more capital flows into these assets than the forced selling via the liquidations. The practicalities of new capital flowing in is worth pondering also. It is possible that this would simply take the form of buying pressure to counteract forced selling pressure; but it is also possible that the buys would be of newly minted assets, hence recycling the liquidity unlocked by collapsing leverage *immediately into new leverage*. This would have a naturally magnified effect on the valuation metrics cited – and shortly to be criticized in subsection iii) – and may help push back against an equally collapsing narrative, should that be a contributing factor.

In one sense, this is all straightforwardly understood as a spiral of debt collapse. If this were all, it would be fair to dismiss our analysis thus far as uninteresting fear mongering given that, clearly, *any* debt can default, and any leveraged financial system therefore bears risk. But this is not our point. It is not interesting that the debt *could* collapse, but *how it could* collapse, and what such a collapse might reveal about its mechanics all along.

There are two factors that are perhaps discordant in this scenario, and that the reader may have noticed: i) what the backstop *really is*, given what the value *really is*, and, ii) how seemingly irrelevant the *cause* of the collapse is to its consequences. While Tolstoy knew well that every unhappy debt collapse was unhappy in its own way, what “normal” deleveragings at least have in common is that expectations of future profitability (i.e. *real yield*) were realised to have been overly optimistic and are adjusted downwards. The adjustment means both the financing costs and fragility of the capital structure erected in more optimistic times can no longer be sustained. But in this case there is no real yield in the first place, so this cannot be neither the trigger, *nor the backstop*.

To tick off i) above, then, the value here is conceptualized by all involved *to be the complexity of rehypothecation, leverage, and securitization created!* What is boasting about the creation of “money legos” if not a celebration of this complexity, regardless of what the complexity either represents or achieves on fundamentals? The backstop clearly cannot be that the underlying *real value* arrives at a more appropriate capital structure once enough debt it could not previously sustain has been washed out because, here, *the debt is the value!* When the debt collapses, so does the value proposition. Crypto, entirely perversely, is most conceptually sustainable *the more debt there is* – which is clearly at odds with its financial stability and, to our minds, has no

resolution. Which all means, of course, that the backstop is fresh external capital. It has always been fresh external capital, and it probably always will be.

As for ii), the trigger, the reader will recall, is simply *prices falling*. If the price in question were a securitization of *rights to the product of all this leverage*, the mechanics might seem comparable, but by far the most likely trigger is a fall in the price of Ether. And if this is not a specific trigger, it is likely at the front of the queue of consequences and hence will quickly *become* a compounding trigger. But what *is Ether itself*? Why does it have value?

This is the root of what we deem to be a severe philosophical discordance. Ether is (allegedly) the right to claim decentralized computation. To make the link more obvious, it is *the right to run a crypto application*. Its price is surely a reflection of “*how much running a crypto application is worth*,” somewhat tautologically? To force the conceptual link to [Section 1](#), then, this *has value* because the decentralized computational resource that contributes to creating the Ethereum timechain is itself scarce, hence this “right” ought to have a market clearing price.

Surely this is something of an operating expense in the scheme of things? Or, at least, surely it *ought to be*? What *should* this have to do with the solvency of the “decentralized financial” enterprise on which it runs? A strict analogy here would be something like a bank that not only has an office and an electricity bill but which also decides to index its leverage ratio to the price it is paying for wholesale electricity or on its rent. And just to rub it in, the analogy demands we rule out the one way this might make some economic sense: we are *not* postulating that the link comes about because rent is a high component of the bank’s costs, hence with profit margins contracting, it is prudent for the bank to prune its own risk. In fact, the exact opposite! “Costs” going up are a good thing and costs going down are a bad thing! Hence a borrower could receive a margin call, not because of anything to do with financing – her own or even *the bank’s* – but *solely* because the bank’s bills went *down*! The perverse justification might be something like: the bank only needs to make so much profit, so with more profit now expected, some outstanding loans can be cancelled.

The reasoning we are forced to follow in this analogy is clearly absurd. But what about the base case that we are using the analogy to try to elucidate? Is it absurd too? To be frank, we think that yes, it is. This is the paradoxical core of all the securitized values pointing at something else and at each other in the absence of real economic productivity. The financial health of the ecosystem depends on the price of an asset that can only conceivably *have value* as a proxy for enthusiasm about the ecosystem given it is a cost that all must pay to participate. But is this really to say anything more than: *the value depends on everybody thinking it ought to have value*? We do not think it is.

Bitcoin arguably fits this description too, but we explained in [Section 1](#) that there is a very good reason for people to believe this in the first place. Or rather: there is good reason for people to believe *other people will believe* Bitcoin has value. Crypto appears not to have this philosophical buttress, and as such, we believe that crypto is incredibly systemically fragile. What if people stop thinking this? What does this perverse mechanism of justifying value churn out at that point? Is there any reason *not to* expect an eventual catastrophic crash, except precisely as perpetually staved off by fresh capital, as repeatedly alluded to above?

New investors daring enough to buy during the bad times would be forced to take a harder look at what “value” they are actually buying into. As much of this piece has touched on, we think when investors cannot simply ride the coat-tails of high beta, the “value” they will see in these assets during this time will be far lower than what is being perceived today. This scenario does not necessarily imply the “end of crypto,” but instead a reset of expectations and beliefs that will bring to the surface important questions about how the ecosystem is structured.

It will be fascinating to see those passionate about crypto cast their arguments during such times. Those that will still have conviction during, and after, a mass liquidation event across the ecosystem are best to be at least listened to and debated, as they are likely the most thoughtful around. For what it is worth, and to foreshadow [Section 5](#), our prediction is that, at such a time, activity will move to then-more-mature higher

layers of Bitcoin. This is, once again, because it is not the *concept* of decentralized finance that we believe is at fault, just *this instantiation*. This instantiation might be thought of as using “*the fact of not having yet defaulted*” as the toxically self-referential basis for leverage rather than real cash flows, real productivity, and a real grounding of value.

iii) misleading metrics

The reader may be aware that, despite our concerns, such headline numbers as “market capitalization” and “total value locked” continue to rise, seemingly indicating ever-growing health and utility of the ecosystem. Is this not contrary to our rhetorical framing that *crypto is not finance*? It seems to be getting more and more financial!

Unfortunately, we believe these metrics are deeply misleading, and we would argue further that the *precise* way in which they are misleading is insidiously what contributed to attracting more fresh capital in the first place. As above, this capital is then rehypothecated, levered, securitized, misleadingly quantified after all this to look wildly more successful than before, and the cycle begins once again.

Our thesis on “crypto valuation,” as briefly as possible, is that none of it makes sense. Similarly to what we will discuss in [Section 4](#), it involves a grave category error that originates in thoughtlessly transferring methodologies from one area of finance to another, without having given the requisite thought to the fundamentals of the methodology in the original environment. This leads to failing to understand *what the methodologies mean and why they mean that* where they do apply, and subsequently misapplying them where they very much do not apply.

Evidencing this abstract objection, we have two specific critiques: i) double counting leverage and rehypothecation in “total value locked,” and, ii) misusing “market capitalization.” The first is easier to grasp and nicely communicable and we borrow an example from Lucas Nuzzi, Antoine Le Calvez, and Kyle Waters’ recent CoinMetrics blog post, [Understanding Total Value Locked](#), the building blocks of which are very much like our own above, albeit more typical and less complex-for-the-sake-of-making-a-point:

“take a look at the following example:

- *A user deposits \$1,500 worth of Wrapped Ether (WETH) into Maker to get a loan in the form of \$1,000 worth of DAI (150% collateralization ratio).*
- *The user then deposits this newly minted DAI, as well as another \$1,000 worth of USDC in the Uniswap V2 USDC-DAI pool. In return, the user gets Liquidity Provider (LP) tokens representing that \$2,000 stake of that pool’s liquidity.*
- *The user can then redeposit these LP tokens into Maker to get another loan of \$1,960 of DAI (102% collateralization ratio).*

From a naive perspective, TVL could be computed as:

Collateral	\$ value
Wrapped Ether backing the original loan	\$1,500
Liquidity added to Uniswap V2 (USDC)	\$1,000
Liquidity added to Uniswap V2 (DAI)	\$1,000
Uniswap DAI/USDC LP tokens backing the new loan	\$2,000
Total Value Locked	\$5,500

Yet, a more sophisticated approach would only count the \$1,500 of Wrapped Ether and \$1,000 of USDC as the “real” collateral giving a TVL of \$2,500. This approach would not include assets that are claims to other collateral such as DAI (which is minted as a loan against collateral), and Uniswap DAI/USDC LP tokens (which represent a claim to the liquidity held by the Uniswap V2 DAI/USDC pair).”

Above we articulated the prospect of rehypothecated collateral as presenting a systemic risk. Here it is clear it also lends itself to misleading quantification of value in the ecosystem. This contributes to a narrative that serves to attract more capital, making this initial problem worse still.

The second objection we have is the importance of the subtle misuse of “market cap.” Market capitalization in equities means the price of one share multiplied by the number of shares outstanding. Although often used as a proxy for *how big a company is*, it has a precise practical meaning: it is a measure of the total capital required to purchase a company outright based on the price at which the shares are trading currently. Usually when one company buys another, it offers a “premium” above the closing price of trading, say 20% or 30%, on the assumption that the majority of voting shareholders will be happy with that instantaneous gain on their investment. That is how sellers conceive of the price of their sale, but the buyers have to buy *all the shares*, hence “market capitalization” has an intrinsic relevance and meaning. Similarly, if a company issues new equity, it will presumably be for the purpose of financing some real-world project (possibly buying a company!) and hence a 5% issuance, for example, has intrinsic relevance in terms of the dollar amount that allows the company to raise. This is why “market cap” is meaningful above and beyond mere “price.”

There is no reason to refer to “market cap” rather than just “price.” And not only is there no reason to; it is arguably actively misleading. Even in regular equities it is potentially misleading because it is an extrapolation based only on the price realised by those willing and able to trade: it reflects the clearing price of the last trade, which may or may not be the clearing price for all shares. It is most often a reasonable extrapolation because in liquid enough markets it is true almost by definition that everybody *is willing and able to trade*, hence if the price really were unfair, or for some reason not reflective of what investors thought the company as a whole ought to be worth, they would act on this intuition and move the market. The reason corporate buyouts are attempted *at a premium* is precisely to account for the portion of holders who are not trading, and hence who cannot be extrapolated over. They aren’t willing to sell at the market price (otherwise they would have sold, given the market is liquid enough) but at 20% higher they surely (mostly) would? Or 30%? The entire point of these negotiations is because the extrapolation has a relevant *meaning*.

But in crypto, and arguably all of crypto, the extrapolation is misleading because *you cannot buy the whole thing*. And so, what if the “market cap” is based on the active, willing, and able trading of 1% of the token supply? Or 0.1%? Or less? Can that *really* mean that there is \$10bn of value “in the ecosystem”? We doubt it. If Alice sells Bob one one-trillionth of his pen for \$1, that does not make Alice a trillionaire. And that

is not to mention the compounding effect of leverage, rehypothecation, and securitization tied up in such metrics as “total value locked” *prior to this point*. We believe that, generally, market caps should be viewed with hesitancy in the space broadly and this is not necessarily crypto specific – financial professionals ought to be just as careful with the idea of Bitcoin having a “market cap” – but in light of the double counting just alluded to, in addition to this theoretical flaw, it should perhaps be viewed *even more* suspiciously here.

This might be little more than a curiosity – or perhaps even us cherry-picking criticisms – were it not for the fact that these numbers are frequently touted as capturing the ecosystem's growth, health, and to some extent, *success*. This in turn attracts the initial capital that we argue is probably necessary to avoid collapse. We will go into much more detail in [Section 4](#) as to the attitudes of institutional investors, *from their point of view*. What we present above is more from the perspective *of the ecosystem* – why does the ecosystem *need* new investors? The reader would do well to remember when she gets to [Section 4](#) that these two perspectives seem to have little to do with one another ...

iv/iv) this is not finance

We want to be clear once again that our purpose in this section is neither to holistically capture the crypto ecosystem, nor to cherry-pick flaws for the sake of it, but rather to draw attention to the many facets and many consequences of one single, dire flaw: that this is not *real* finance. Its foundation is fundamentally economically suspect as it can be rooted in neither spot appraisals nor yields. There can be no spot appraisals, akin to Bitcoin, because the key innovation that allows Bitcoin to realize emergent and endogenous value has been removed in the process of creating these assets. There can be no yields because the “returns” being generated and oft-cited derive from leverage, rehypothecation, and securitization, not links to productive capital. What value undeniably exists is the derivative of primary capital that is “deployed” only back into more complex manipulations of itself, not intermediated to real investment, hence not finance.

A financial ecosystem built on products with value propositions (and in the stablecoin cases, *actual values*) that are not a reference to stocks of productive capital and the flows they generate, but rather point back and forth to each other, cannot be justified as some kind of brilliant new design. It is just leverage on leverage, rehypothecation on rehypothecation, and securitization on securitization. Sometimes the leveraged assets are rehypothecated, and sometimes the rehypothecated assets are leveraged. Sometimes the claims to products of leveraged and rehypothecated assets are securitized, and sometimes the securitizations are used as “collateral” for more leverage. But this obscurantist complexity is a bug, not a feature. This is compounded by the fact that the “real world value” backing any of these assets – *even the initial collateral* – is itself unclear.

Lyn Alden captured this well in her excellent piece, [An Economic Analysis of Ethereum](#), which we highly recommend in addition to our own. Similarly observing that there does *seem to be* a lot of “financial” activity on Ethereum, specifically: it’s what you would build if you were trying to replicate finance just by looking at it but not really understanding what the *point* of it all was:

“Ethereum is heavily used for decentralized exchanges of crypto tokens, crypto stablecoins that serve as liquid units of account for trading crypto tokens, and lending and earning interest on crypto tokens which is a practice that serves as a liquidity/borrowing source for traders of crypto tokens. To a lesser extent, it is also used for gamified ways to earn or trade various crypto tokens.

So, it’s a big operating system powered by crypto tokens, for the purpose of moving around... crypto tokens.

A healthy banking system in the real world would consist of people depositing money, and the banks making various loans for mortgages and for business financing, to generate real-world utility.

A speculation-based banking system, on the other hand, would consist of a bunch of banks taking deposit money, and then lending to speculators in the nearby stock market, along with technology providers that make this easier, and then what those speculators are trading mostly consists of shares of those banks,

shares of those tech companies, and shares of the stock exchange, resulting in a big circular speculative party. The biggest use case so far for Ethereum is a decentralized version of that circular speculation-based system.”

Our only disagreement is that Ethereum is not properly “decentralized,” but we covered that already ...

The reader might wonder if all of this criticism might not equally apply to Bitcoin? There are multiple parts to this answer, some of which we have implicitly addressed already, and some of which we will address later in the piece. The simplest retort would be that the base layer of Bitcoin is only trying to be money, and nothing else. We will explain in much more detail in [Section 5](#) that we anticipate similar “crypto” tools and ecosystems will come to exist on Bitcoin, but that precisely this grounding in *real value* is what will make them more secure. We might go even further – and in doing so, further buttress our philosophical defence of the innovation of proof-of-work, the difficulty adjustment, and the truly distributed consensus they enable – by arguing that “money” is an endogenous and emergent use case of the Bitcoin timechain, and that, actually, this emergent value is ultimately dependent on whether users are happy enough with the security expense to continue paying it. This, of course, gives us our “real world value” link: censorship-resistant, integrity-assured floating value native to a distributed ledger is well worth paying for with time and energy, hence people do. In Bitcoin, everything ultimately points to proof-of-work.

If attempting to follow the *real life* analysis and explanation above unfortunately proved too confusing, we can alternatively model a hypothetical altcoin valuation mathematically as follows: Let n be a natural number denoting discrete steps of *time* in the model. Then let i, j , and k be natural numbers referring to crypto tokens which index both $DF_i(n)$ and $VC_j(n)$ as follows: $DF_i(n)$ refers to a valuation for crypto i -token at time step n as facilitated by a primary equity raise to support the asset price in the open market. $VC_j(n)$ refers to the flow of capital to the venture capitalist, either primary or secondary, and either positive or negative, based on whether capital is being injected into the market or removed from the market as a “return,” respectively.

Keeping in mind that much of crypto “valuations” are “justified” on the basis of value locked in other crypto assets they nominally either custody and rehypothecate or over which they claim some “governance rights,” (i.e. *of which they are securitizations*) consider the following (not entirely serious) model:

$DF_{j(n+1)}$: a primary equity raise facilitating a secondary transaction $VC_{j(n+1)}$, on a valuation uplift justified by rehypothecating/(re)levering custodied / securitized “governed” assets whose paper value is based on the market price as determined by the open market trading following $DF_{k(n)}$, for $j \neq k$.

with:

$$\forall i: VC_i(0) = DF_i(0) = \text{value of } i\text{-token at launch} > 0$$

$$\forall i: \forall n \geq 1 : VC_i(n) < 0$$

$$\text{and, } \forall i: \forall n \geq 1 : (1 + (-VC_i(n) / VC_i(0))) ^ (1/\text{years}) - 1 = IRR(VC_j) = \sim \text{one bajillion } \%$$

(graphical representation overleaf)



Figure 2: Cartman LARPing as a “Crypto VC”

As a final metaphor that is easier to understand than a barrage of quasi-satirical mathematical logic, and which we think is still accurate enough, as far as it goes, and hopefully helpful, consider the following model of crypto funding:

Venture capitalists build a Jenga tower, contributing 100 Jenga blocks. They are invested to the tune of 100. Then we get into an iterative game in which the VCs take 1 block from wherever they like, while retail adds 2 exclusively to the top of the tower. By the time VCs have taken, say, 20 blocks from throughout the structure, they have a paper return of 20%. Of course, it is “paper” only because, unlike a real yield from a real asset, there is absolutely no reason to believe the par value of the principal is assured. Eventually, when the “returns” are impressive enough to justify it, VCs raise much more primary capital, which, this time, *does go to* the bottom to buttress the structure. This means the game can go on for longer, but it doesn’t really do much more than that.

The point of Jenga is *not be the one* who collapses the tower – because it *will* eventually collapse. At the risk of yet more mixed metaphors, this is all rather like a game of musical chairs: there mathematically is not enough for everybody to make whole, never mind make a “return,” but, as long as the music is still playing, nobody seems to care. We cannot help but feel that the crypto game, as of today, is ever-so-slightly more complicated: to recoup a non-paper return by extracting more Jenga blocks than you contributed ... before eventually the music stops and the tower collapses. To be clear, none of these “towers” have collapsed in the way we are alluding to now, but we suspect that this is because capital has continued to flow in, and that, if or when this stops, collapses are likely to follow.

While some of this may come across as unserious, we have done our best to be accurate such that the reader is getting a sense of the typical construction in the space. The value involved is clearly speculative – even proponents would surely not deny this – but there appears to have been curiously little thought put towards *why* speculative value ever comes about in the first place. In this case, what is being speculated on is that something will cease to be speculative and become real. The problem is that these are the same thing:

it is a speculation that this speculation will stop being speculative.

Without the tie to real world assets, and simply having fresh capital flowing in as a backstop, we will remain seriously concerned about the true health of this ecosystem. And of course, as discussed in [Section 2](#), any tie to real world assets we strongly suspect is an unacceptably risky and borderline pointlessly expensive misuse of this technology over what can be achieved with regular computational tools.

We can push this even further in tying the discussion back to Sections 1 and 2 and setting up the link to [Section 4](#): we believe the speculation is tied up in *yet another* vicious circle. The assets themselves need to be expensive in order to thwart attack by raising the escrow stake and the dishonesty penalty for a would-be attacker. They need to *do something productive* in order to justify their expense beyond mere speculation (or, as above, the speculation must be based on the reasonable assumption of *one day* doing something productive). But they also need a near-enough guarantee they will thwart any attack in order to be reasonably speculated upon as likely to one day do something.

Bitcoin must also prove its resilience to this cycle – but for Bitcoin the argument is not only straightforward but fundamental to understanding the *only* thing it does and does *well*: what it *does* is provide an appraisable spot value (not a *yield*) as censorship-resistant, integrity-assured floating value native to a distributed ledger. Therefore, expense is justified, therefore attack is thwarted, therefore there is good reason to speculate on future floating value.⁵ In Bitcoin’s case, it is really more of a *virtuous* cycle, in *virtue* of the genius of proof-of-work, the difficulty adjustment, and the censorship-resistant and integrity assured distributed consensus they enable.

If crypto is likely not resistant because it is not decentralized, and not productive because of financialization rather than finance, then what is the basis of the speculation propping it up? In [Section 4](#) we attempt to address this. Surely the speculators do not conceive of their capital allocation as, *speculation that this speculation will stop being speculative*. So, what do they think? What is the investment rationale?

⁵ There is a reasonable debate to be had about exactly how secure Bitcoin is and exactly what the threats are, but it is outside the scope of this paper. Our own view is that while skepticism is healthy, we do not in general worry about Bitcoin’s long-term security. But we have patently given no such argument here and readers are strongly encouraged to do their own research.

4/n – The “Investment” Rationale

“More succinct advice to those who must time markets comes from remarks attributed to a nineteenth-century cotton trader: ‘Some think it will go up. Some think it will go down. I do, too. Whatever you do will be wrong. Act at once.’”

David Swensen, on “Crypto”

TLDR: *in Section 4, we provide a rationale for investing in the space to date and argue that the most sensible investment thesis is a subtle category error that results in transferring across principles from software venture investing that do not **quite** apply in this space. We argue further that only a certain class of investors is likely to commit this error, and the realization that others will not follow will likely mark the beginning of the end.*

Jump to [Section 5](#) for the argument that precisely the desirable features of crypto will likely emerge before too long on Bitcoin.

This is not a psychologization or pathologization of crypto proponents. The “investment” rationale is a category error that arises from misapplying the traditional venture capital methodology in an area similar enough to seem familiar, natural, and possibly even identical, but *just different enough* to be deeply philosophically questionable. We will get to this later on but, of course, we could be wrong; perhaps we ourselves suffer from a pathology of “maximalism” toward the likes of: sound money, layered architecture, provable technical robustness, long-termism, and capital formation, that blinds us to what is nonetheless worthwhile innovation, *and* an abhorrence of mainstream Cantillionaire finance leading to a preference of prospects for its replacement rather than its being made even more digital, shorter-termist, and “efficient.” It is without a doubt the case that many intelligent people dedicate their time to crypto and will disagree. Their rebuttals are welcome and we hope this piece will prompt them in good faith.

To set the stage for an analysis of what we believe is the “category error,” the following is a summary of the methodology and rationale for investing in the equity of ultra-high-growth, early-stage software companies:

Software has a unique economic characteristic of being unboundedly reproducible at near-zero marginal cost. Companies producing and looking to profit from proprietary software are amongst the most inherently operationally leveraged in the history of industrial capitalism. This creates strategic priorities for early-stage software investors that the best venture capitalists figured out a long time ago, and that might look crazy from a traditional capital allocation and business development perspective without having realised these fundamental economic differences.

Software businesses typically have enormous addressable markets, at least in principle. Even the most niche application imaginable (let’s say, SaaS for SaaS companies serving industry niche X to manage their industry niche X SaaS subscriptions) can theoretically address every such use case in the entire world from day one given that the customer requires only an Internet connection and payment to the merchant requires only the kind of infrastructure offered by Stripe, Adyen, Square, PayPal, etc. that is nowadays ubiquitous. Moving to a new jurisdiction or even a new country may require next to no operational infrastructure. Without meaning to sound flippant, there is likely already SaaS for that. It is just another operating expense. Therefore, it makes eminent financial sense to grow as fast as can be done without incurring the risks of truly excessive operational, or potentially *financial*, leverage as brought about by growth strains, runaway costs, no profitability, cultural dilution, and so on.

This is often misunderstood as requiring or even indicating “network effects.” This can be true but is unnecessary and only peripherally related to this analysis. The presence of potential network effects in a product or service simply makes this issue all the more pressing but they are not necessary. What matters is that this economic profile heightens competitive pressures to levels unimaginable in any “normal” industry – i.e. one without the bizarre and relatively novel underlying economic profile. As a software start-up, you have to capture the market as fast as possible for purely game theoretic reasons: if you don’t, somebody else will. Perhaps more importantly, *everybody else can*.

One can’t point to the return on capital implicit in your unit economics and say, *but it’s unsustainable to try to grow faster than that!* Wrong. What’s “unsustainable” is losing the race to win the market and going out of business entirely. What’s sustainable is to cover operating losses brought about by rapid growth with financing until the market is all but won. It is sensible to treat all this as, effectively, R&D, although impossible under accruals accounting standards developed a good 60 years or so before software as we know it today existed. Some of it may be *actual R&D* – i.e. in *technology* – but in spirit, it is R&D into company design. The company is running profitless experiments aimed at discovering what it ought to one day look like when profitable.

There is a final, added wrinkle. Even this may not be sustainable for an individual company. Early stage software companies *in isolation* are amongst the riskiest and most uncertain investments that can be made. But those that succeed are amongst the best investments that can be made. This is why venture capitalists have portfolios. Not because *there are lots of good investments out there*. The single best portfolio is the single best investment. But nobody knows what the best investment is going to be and VCs have a fiduciary responsibility to their LPs who don’t care about the success of individual companies but about the return on their entirely impersonal and illiquid investment. VCs offer a service to their institutional clients that may be roughly described as: returns both above and uncorrelated with equity and bond market indices, the (hoped-for) premium of which is worth the comparative and required illiquidity. There is a fascinating debate to be had about the optimal VC portfolio size, given too wide a spread dilutes the impact of outliers, but too narrow a spread risks missing winners entirely.⁶ Regardless, the minimum cut-off advocated for by any sane VC is certainly above 1. The minimum the authors can recall seeing justified in terms of this debate is around 12-15 and the maximum is effectively unbounded as it becomes more a question of operational bandwidth than purely financial reasoning.

Going into such detail on the financing logic is necessary as it affects the operational logic. VC software equity investing is only justified on the basis of its contribution to portfolio level return on capital. This aggregate return dynamic is in turn predicated on the realization that most will go to zero, but portfolio returns may still be good overall based on one or two dramatic outliers. Hence there is no point in an individual company trying to be merely “good.” Everybody must endeavour to be great, few will succeed, most will be awful, and the average will wash out as “good,” which was the point all along.

This operational logic wraps back around to affect financing logic also. Slightly rephrasing the description of early stage software businesses for optimal suggestiveness: what it *means* to “*cover operating losses brought about by rapid growth with financing until the market is all but won,*” is to willingly sacrifice the ability to generate a *yield* for a rationale that might seem crazy at the level of an individual company but makes perfect sense at the level of a portfolio, so as to maximize the *potential* yield the individual company might one day be able to generate. The financing provided is an investment in developing the economic carrying capacity of the firm. It clearly takes time. Creation of real productive capital always does. It is speculative, but it is obviously sensible, healthy, and an enormous positive sum for the world at large. Venture capital is financial engineering to funnel low-risk savings to extremely-high-risk investment projects. The net and aggregated capital creation is nothing short of phenomenal.

But ... this does not apply in crypto. Even worse, all of it *seems like it applies*, if you haven’t really thought it through. To start with, we do not think there is a likely final, profitable, sustainable, *yielding* state to

⁶ Jerry Neumann’s wonderful *Reaction Wheel* blog covered this [here in 2015](#), and again [here in 2017](#).

be aspired to, as was exhaustively described in [Section 2](#). We think there is also no realistic prospect of this ever existing, as explained in [Section 3](#). This might be countered, as follows:

“Sure, it’s a low probability of success and many will fail, but this is innovative R&D, so it makes sense to diversify across a portfolio in the interest of long-term ecosystem support. We are building alternative financial infrastructure that, once proven, can be plugged into the real economy. But in the meantime, the only way to test the underlying theses is to provide upfront liquidity and see what sticks.”

No, no, no, no, and no.

Given our concerns, we put it out to the reader that the probability of “success” is very, very low. We cannot claim to *know for sure*, but we suspect all will fail for fundamental philosophical, technical, and economic reasons, on a long-enough time horizon.

It’s not the kind of R&D that is rational in these circumstances. It is purely *technological* R&D, which may be perfectly commendable out of context, but this context is crucially not one of charity. Very little of this R&D is going towards the all-important discovery of what a profitable end-state might look like. So, the funds will eventually run dry once fresh external capital no longer enters, and the R&D will, by necessity, cease.

It doesn’t make sense to diversify because the point of diversification is *not* that there are lots of interesting opportunities but to try to wrestle the space of portfolio-level returns to something more manageable than “wild uncertainty” and, in doing so, to satisfy fiduciary responsibilities to LPs. This logic in early-stage software businesses depends on small probabilities of enormous payoffs. But, as just covered, we believe the probabilities here are effectively zero, so this logic falls apart.

Applying this thinking to crypto is not long-termist. It is frankly much more similar to blowing up a bubble. Of course, long-termism can certainly be speculative. The financial engineering of *real* venture capital, as just analyzed, is amongst the most long-termist enterprises in contemporary finance. Given the time periods over which transiently unprofitable, and potentially never-to-be profitable, projects must be patiently supported, and which would probably never otherwise come into existence. But this is all predicated on the creation of economic carrying capacity and the *potential* for a yield, one day.

In crypto, the speculation seems to us to be entirely self-referential and so the narrative has become alarmingly *short-termist*. The narratives seem to constantly change and yet immediately be very well publicised around buzzwords nobody has ever previously heard of. The desired impression seems to be of innovation that is spiralling out of control, even though nobody seems bothered by the fact that the buzzwords and narratives from two such cycles ago came to precisely nothing:




Figure 3: fads in crypto marketing over time

This affects financing at both ends. It generates retail interest such that secondary trading drives up mark-to-market prices and validates “returns,” *and then* those “returns” are used to justify raising more primary capital from FOMOing LPs.

We say “returns” rather than returns, because none of this is real. There is no yield currently and we see no path to a yield, either. Which is why, finally, **it won’t ever be plugged into the real economy, either transparently or at scale**. Real institutional capital allocators have an entirely different kind of responsibility towards their clients than the VCs currently engaged in this space. The speculative element *currently* is an investment risk, not an operational risk. The reader might be tempted to argue that the whole point of *decentralized finance* is that we no longer need institutions because they are being automated away. This might sound nice but it is pure fantasy. As per [Section 2](#), retail is only involved with secondary trading driving mark-to-market valuation hikes. 99% of the primary capital is institutional because these secondary trades are providing it with incredible paper returns.

The secondary trades, hence the returns, depend on liquidity, and liquidity depends on permanence, in a way that is so obvious to never usually need to be spelled out like this. If market participants *know* a security is worth zero, then *liquidity* will only continue to exist, so long as there is belief in a greater fool to sell to. In this case, that greater fool is retail. Or, put another way, we must not mistake market liquidity and market *depth*. Perfectly deep markets are liquid, right up until everybody wants to sell.

This upfront liquidity is highly ill-advised, because it is facilitating *only more transient liquidity*. A massive amount of capital and a large number of “investors” are involved, not for anything like the reasons outlined above, but simply because this liquidity transiently exists and seems like a worthwhile way to “diversify” capital at the portfolio level.

To stress the dire state of transparency and liquidity in this ecosystem even further, we also happen to know of large volumes of SME loans currently being used as collateral tranches for stablecoins which we are fairly certain is *not* well understood by either side of this transaction. This is why we said “transparently” above: there *is a connection* to the real economy here, but it is not transparent in the slightest. Crypto is serving as an extension of shadow banking. These are unregulated eurodollars in all but name.

This is to highlight *yet another risk*, and to be clear on what is meant by “links to the real economy.” We do not believe these count towards the loftier goal of “links to the real economy,” because they are really yet another form of ignorant and desperate yield chasing and comfort in diversified liquidity, regardless of solvency, *on either side of the trade!*

Next, it is worth briefly drawing attention to the incentives of current market participants. The incentives of retail and institutional *end clients* deploying capital in this space is *chasing yield*, because every other asset class is inflated beyond all reason by central bank intervention in financial markets. But this is not how VCs are paid. Their skin in the game is not a stock, but the promise of skimming flows of this stock. They benefit directly from however long they can continue to deliver “returns” on which they are entitled to skim 2 and 20, and only indirectly from the solvency of the paid-in capital. If the paid-in capital *looks* solvent, the skimming continues. It doesn’t really matter whether or not it *is* solvent.

This might seem harsh, but we can testify to the existence of novel services industries sprouting and skimming non-yield-flows from this capital consumptive space, even beyond the relatively more tangential anecdotes above. In recently speaking to management of a public company considering “launching a blockchain,” being as vague as possible so as to keep them effectively anonymous, we were informed that the marketing strategy was to hire a PR firm to manage a worldwide network of “YouTube, TikTok, Clubhouse, and Twitter *crypto influencers*,” to encourage retail investors to “support the token price upon launch.”

This is a shadow industry. It’s not entirely unlike intermediary subprime mortgage boosters, in our view. It is certainly at least as questionably legal. Given this is a philosophical, technical, and economic analysis, this should not be misunderstood as agitating for prosecution for unregistered securities issuance or promotion, or regulatory intervention of any kind. Our point is simply that, regardless of whether or not this *should* happen,

[it will](#). Given crypto projects are almost all de facto centralized to a greater or lesser extent, as explained in [Section 2](#), such prosecution would likely *work*.

When Joe Weisenthal and Tracey Alloway said the following in their introduction to [interviewing Alyse Killeen on their *Odd Lots* podcast](#), we couldn't help but chuckle:

Joe: *"I feel like, in some sense, some of the narrative enthusiasm has moved away from Bitcoin. At least it feels like it in the last few months."*

Tracey: *"Oh, absolutely. I think part of this is because obviously you and I are in financial journalism and we are talking with people who are in the financial industry quite a lot. But it feels like the enthusiasm from traditional financial players – you know, I'm thinking bankers and traders – that has squarely moved on to DeFi, to things like Ethereum, to places where there seems to be a lot of innovation and changes happening around what you can actually do with the technology and with a wider pool of crypto."*

It is no surprise that those whose professions and entire capitalistic purpose is to skim short-term flows without contributing to the long-term creation of real, productive capital are increasingly attracted to a space in which yieldless flows are flying around for whoever wants to grab them. On top of this, those that are losing interest in the alternative space, in which productive carrying capacity is being created, will one day make these very same people's employers and professions obsolete.

We are shocked! Shocked! Notice that certainly Weisenthal and Alloway, but probably the bankers and traders to some extent too, seem to have fallen for the idea that "the narrative" is organic, rather than a sophisticated media operation – or are at least a little too open to entertaining this idea for our liking. These are the same kinds of people who say things like "*trading is a use case*," and who think "liquidity" is an end rather than a means.

Ammori makes a comment in a similar vein, even more confidently, writing:

"A by-product of building financial services on a transparent, shared database is that all associated transaction data is publicly available in real time. For example, earnings generated by liquidity providers in the Uniswap Protocol can be tracked on per-second granularity. Investors can use this data to decide how to allocate capital, providing for more efficient price discovery and allocation of resources, while regulators can monitor real-time transaction data to identify nefarious user activity.

This is a significant departure from traditional capital markets, where investors are left entirely in the dark until firms issue their quarterly earnings reports. The state of private markets is even more dire, with companies often inventing their own accounting metrics, if they decide to release metrics at all. It is difficult to imagine investors making rational decisions when having to work with stale data!"

Ammori is transparently *not* talking about *investors* in any coherent sense of the word, but *speculators*. Not *finance* but *financialization*. There is nothing remotely economically useful to be gained from "real time financial data," and much to be socially lost given it encourages yet more time be wasted on the metagame of "**allocating** capital" rather than the game of **creating** capital. Investors *should* be left in the dark because otherwise business operators will be spending all their time briefing ignorant financiers rather than operating businesses. If anything, quarterly reporting is far too regular as it is. No worthwhile capital formation happens over three months, and the idea that it does – or *can* – is [cargo cult math](#), to which financiers are typically highly susceptible, but to which operators tend to have greater natural immunity.

The ultimate irony is that Bitcoin fixes this. Not just "crypto" non-use cases but the fundamental reason this capital bonfire exists in the first place, central bank intervention in financial markets. The capital being deployed in this space is *chasing yield* for reasons already mentioned. There is a perverse portfolio-level concern at play here that oddly mirrors that faced by venture capitalists. The headline numbers highlighted in [Section 2](#) may all look enormous, but these are mere basis points for the end clients. This is their funny money. The biggest US public pension plans and family offices are the ones ultimately on the hook for all of this, as well as

whatever retail gets swept up in the hype. They don't have the time to do the kind of work we are doing right now because they invest at least a little in every single investment opportunity in the world.

But that is not to say they are not sophisticated. They are amongst the most sophisticated investors around, or at least the most sophisticated *capital allocators*, precisely on the basis of their high-level view across the comparative merits of different asset classes. They can afford to put basis points into this experiment because they do not believe, as we do, that no real yield will be generated in the long term.

We opened this section by noting that there is no intent to psychologize or pathologize. But we will conclude by breaking our own promise – *positively!* One cannot *know*, but can strongly *suspect*, that the personal motivation of those involved with good and noble intentions are very likely some variation of, *we need to rebuild the Internet, and, in doing so, have a shot at fixing finance as well*. Pressed further, they will likely give entirely accurate and laudable critiques of the way the Internet and the financial services industry are currently designed. Nonetheless, that their diagnosis is sound tells us nothing about their proposed solution.

As Dhruv Bansal and Ryan Gentry argued in [their talk at Bitcoin Miami](#), the broken architecture of the Internet *is ultimately caused by central banking*, and the contemporary design and functioning of money. Or, a little more specifically and less anachronistically, the lack of a digitally native censorship-resistant and integrity-assured programmable money. Bitcoin fixes this. Altcoins do not because in trading censorship-resistance and integrity-assurance for *even more programmability*, they break the innovation and will likely themselves be broken as a result.

In case it is unclear, we aren't attributing any malice to those doing this work. We know many personally and wish them all the best. Many were kind enough to contribute feedback to this piece. If we are wrong, they will be fine. If we are right, we echo Dhruv's comments in his and Ryan's talk linked to above: we hope they bring their projects – and more importantly their *talents* – **to Bitcoin**. This will have a far greater chance at making a real long-term impact, as explained in the following section.

5/n – Layered Architecture And Gall’s Law

“A complex system that works is invariably found to have evolved from a simple system that worked. The inverse proposition also appears to be true: A complex system designed from scratch never works and cannot be made to work. You have to start over, beginning with a working simple system.”

John Gall on “Crypto”

TLDR: in Section 5, we argue that the desirable features of crypto DeFi will likely emerge before too long on Bitcoin. Furthermore, we argue that the fact of these features taking longer and being more difficult to build is a fundamentally **good thing**; it reflects that Bitcoin’s architecture has been built in a more robust and prudent manner than its crypto peers. Ironically, in the long run, this is likely precisely what will enable extension of functionality. We give basic details on a handful of relevant projects before analysing this dichotomy in more philosophical detail.

Jump to [Section 6](#) for a list of the most obvious reasons our entire thesis might be wrong.

We believe the alleged promise of crypto will be fulfilled on Bitcoin sooner or later. The necessary development on Bitcoin is (literally) years behind its crypto peers, but once again we would argue this doesn’t matter and, in fact, is a *good thing*. The contrary view is the category error of misapplying rules of thumb of ultra-high-growth, early-stage software equity investment and business development. Bitcoin does not move fast and break things. It moves slowly and breaks nothing. If, or when, crypto breaks, Bitcoin will probably still be going strong, slowly and steadily winning the race.

This section will first highlight a handful of projects that envisage bringing greater functionality to the Bitcoin space, rooted in the censorship-resistant, integrity-assured, floating value native to a distributed ledger that is the Bitcoin timechain. This is not an endorsement of these projects or even a prediction of their success. Many, or all, may fail. Very limited information is included on each, more to pique the interest of the potentially curious reader than to really *explain*. But we will secondly go into more detail on why we think these projects, in spirit, embody a superior engineering ethic to that of crypto: layered architecture. In fact, while not hoped for, it would be interesting and perversely beneficial if at least one of these high-profile projects *did fail*, purely so as to then have it be totally obvious that Bitcoin itself is unaffected.

Lightning: is probably the best known “higher layer” of Bitcoin given its obvious and growing potential for internet-native payments. Lightning uses the engineering construct of a “payment channel,” in which a smart contract on the timechain is configured such that two participants can be thought of as escrowing Bitcoin and trading offchain receipts with one another instead of committing to an onchain transaction. These “receipts” have two clever properties: first, either party is able to use the most up-to-date receipt to “close” the channel by triggering an onchain transaction that settles the outstanding balance. This cryptographic enforcement mechanism is what makes the offchain trading of “receipts” effectively indistinguishable from transferring “real” (i.e. onchain) Bitcoin to the extent the Lightning Network itself remains functional. Lightning can hence be thought of as *extending the functionality* of the mainchain to the properties of the network such as instant (effective) settlement and very low cost. Second, the receipts are such that they can be relayed across more than one channel and simultaneously settled across all channels involved: if Alice has a channel with Bob, and Bob with Carol, and Alice wants to send Bitcoin to Carol, there is a similar mechanism of cryptographic enforcement by which Alice can pay Bob *if and only if* Bob pays Carol. Bob has no reason *not to* pay Carol given the enforcement mechanism means he knows he can both trigger Alice’s payment *and* get a very, very small fee for his service.

Lightning is clearly useful in terms of this basic payment routing and settlement functionality, but is arguably far more exciting with a more philosophical perspective: we can think of Lightning “payments” as consisting of data relaying encrypted and onion-routed payments instructions, but we can also invert this interpretation. We can instead think of them as *payments for encrypted and onion-routed data!* This may well be far more profound in the long-run as it potentially solves the classic problem in distributed systems of who is motivated to subsidize the means of achieving privacy. The largest such network on the public Internet, *TOR*, has struggled to scale beyond hardcore privacy enthusiasts given the added cost of operating a node, and, perversely, marks out precisely those who want to act privately or facilitate privacy for others. Lightning, on the other hand, embraces privacy not only as a default but as more efficient than any alternative method of data transfer method *given it is directly self-monetized.*

This realization lets us expand our understanding well beyond *just* privacy and *just* payments. Privacy is a use case worth paying for, and better still worth making the default by a form of coincidental universal subsidization. But *data transfer that is better monetized* is a more general use case, for which privacy is surely a bonus. We believe that the ability to build monetary transfer directly into data transfer will prove revolutionary for the architecture of the Internet. It has already enabled the likes of Sphinx Chat, a chat and content streaming app that can be thought of as “third layer” in that its data transfer between instances *literally is* Lightning payments. But we strongly predict that, on a long-enough time horizon, more or less *all* online data transfer will have a Lightning component, from the likes of metered API calls on the more technical end to pseudonymous identity verification and streamed content monetization on the more consumer facing side – all because it would be foolish not to. It would be throwing away revolutionary functionality.

Speaking of revolutionary functionality, we will close this blurb with the observation that the the taproot protocol upgrade expected to go live on the network in November brings the tantalizing prospect of native and programmable asset issuance. Although we caution the thinking on this possibility is very early stage, if successful, we can’t help but feel it would spell the end of any credible “use case” for just about every crypto token. You very much will not *need a token for that.*

Liquid: enables the extension of Bitcoin functionality using the engineering construct of a “sidechain,” adding such features as asset issuance, confidential transactions, and more. The rough idea of a sidechain is to consciously trade off speed and programmability against decentralization and trustlessness (as opposed to assuming that all can be built into the same “blockchain,” for example, with no necessary trade-offs or adverse consequences). Hence Liquid is deliberately more centralized than the Bitcoin timechain (and the Lightning Network as just mentioned, for context) but is intended to be used in environments in which this doesn’t matter or is expected, while retaining security and robustness as “borrowed” from Bitcoin itself. Liquid has 1 minute block times, with finality after 2 confirmations, and a fully expressive scripting language for programming smart contracts which manipulate Liquid Bitcoin (L-BTC), the asset native to the sidechain ecosystem. Bitcoin is pegged into the sidechain network to unlock Liquid Bitcoin (L-BTC) so there is always a 1:1 relationship between BTC and L-BTC. Asset issuance in Liquid can be used to create stablecoins, game currencies, NFTs, and security tokens. It is also possible to create Lightning Networks on top of Liquid Assets, such as USDT, which would enable stablecoins to be used as payment methods.

Discreet Log Contracts: provide Bitcoin a way to implement any kind of arbitrarily-defined contract directly on the base layer of Bitcoin. DLCs use an oracle or set of oracles to attest to an event that happened outside of the bitcoin timechain’s own data, and this attestation is used to execute a contract. These oracles can attest to anything from who won the Super Bowl, the bitcoin price, or the movement of the stars. This potentially provides the means to build not only powerful contract structures, but staples of crypto, such as [stablecoins](#). DLCs are also very basic and flexible in their structure so they can be used in conjunction on almost any of Bitcoin’s higher layers. There are already proposals for how to implement DLCs on sidechains and even routed across the Lightning Network.

RGB: is an attempt to keep relevant data about the validity of asset transfers and other financial contracts outside the relatively scarce, expensive, and public (i.e. *not private*) global consensus of the Bitcoin timechain, and instead validate such data in an off-chain, client-side setting. The thinking behind RGB is that the mining process only really needs to include valid transactions in the timechain, and hence there may be a

number of benefits if there were a way to move the burden of validation to only those nodes involved in a truly peer-to-peer transfer. While ways to achieve this for bitcoin the asset itself are still speculative, Bitcoin-based non-bitcoin assets are the perfect candidates with which to experiment and trial such a design. They would be *client-side validated* by definition, since miners can't and won't validate the rules of non-bitcoin asset schema. There is also hope that RGB can provide building blocks for other off-chain, higher layer projects to plug and play to ease the burden of their own native technical hurdles (or potentially overcome them entirely) such as always-online requirements for nodes or users, a bloat of data that must be retained but cannot be shared, and so on.

Miscellaneous: the space is proliferating faster than we can claim to understand or keep track of, but to flag just a few more higher-order layer applications: each of Lightning and Liquid are starting to see higher order application layers, such as Sphinx Chat and Impervious.ai in the former case, and Hodlhodl and Bitmatrix in the latter case. But there are yet more ways of creating higher layers on the basis of different forms of integration with the Bitcoin timechain.

These are typically aimed at capturing some notion of “smart contracting” and creating a range of developer tools for *truly* decentralized finance, mindful of the precise security tradeoffs, given *how* they interact with Bitcoin. This is rather than going for everything, all at once, as is seemingly the norm in crypto. Given the limitations of Bitcoin's scripting language, and the importance of these limitations as having emerged from security tradeoffs, there are, roughly speaking, two options for tying the state of smart contracts to the Bitcoin timechain: i) the smart contract states are off-chain and require some or other clever design of offline coordination, or, ii) a separate layer that uses a native asset representing a right to pooled computation, an equivalent of Ethereum's “gas,” to execute contracts that have global networked state.

RGB and DLCs arguably fit option i), as does a smart contracting platform called RSK, [one fascinating characteristic of which](#) is that its smart contracting engine was originally forked from the Ethereum Virtual Machine. This has allowed it to be compatible with Ethereum-based smart contracts. One wonders in that case why a developer wouldn't opt to run the exact same application for *decentralized finance* in a system with dramatically better security on the basis of philosophical, technical, and economic coherence. Unfortunately, the answer is very likely that computation in RSK is “paid for” with R-BTC, pegged 1-1 with Bitcoin, in the vein of L-BTC on Liquid: i.e. there is no potential for a speculative frenzy of token appreciation. Stacks, on the other hand, explores the design space of ii) and aims to bring Ethereum-like generalized smart contracts to the Bitcoin ecosystem without sacrificing Bitcoin's security guarantees.⁷ Stacks is not without *technical* controversy, as we note in the footnote below, but it is additionally interesting for what we can only think to call *social* reasons given the project initially started out trying to build all transactions into Bitcoin's timechain. To paraphrase what its creator, Muneeb Ali, told us (the authors) in private communication, the Blockstack team realized exactly the thesis we are proposing in this document many years ago and moved the smart contracting to a separate layer. The curious reader is directed to [an excellent writeup of the history and attributes of Bitcoin sidechains](#), written by Sergio Demian Lerner, the designer of RSK, which discusses Liquid, RSK, and Stacks.

Projects of this sort are expected to continue to proliferate, be they higher layers on Lightning or Liquid, *higher layers on higher layers on Lightning* (as seems imminent whenever a functioning app is built on

⁷ Stacks is controversial within the Bitcoin community, as its native token, STX, is seen by some as implicitly competing with Bitcoin for a store of value, hence, the fight for liquidity, and, with reference to crypto, as a red flag *in general*. What most importantly follows our own analysis above, however, is that it seems like STX *does not* fall into the vicious circle of unjustifiable value as criticized above as the token plays no part in the *security* of its own ecosystem, although clearly it has other relevance, and it is desirable for the community for the token to appreciate. But Stacks' security relies on Bitcoin's security, so the primary logical focus of our overall critique of crypto seems to have been avoided. STX is only trying to be “gas,” and not to dichotomously double as computational credit and security incentive. Overall, the project seems to us to be an honest attempt to build on Bitcoin that, if it fails – whether on account of potentially misaligned incentives created by the native token or otherwise – poses no risk to other such projects, and if it succeeds will likely provide benefits.

Impervious.ai, for example) or the discovery and experimentation of more novel techniques for building directly on Bitcoin, such as RGB, RSK, and Stacks that will inevitably continue.

Many, if not all, of the above are often lazily described as clunky workarounds to the technical limitations of the Bitcoin timechain. But we vigorously reject this notion on the technical grounds that layered architecture is objectively optimal engineering. Cramming all the features of Lightning, Liquid, DLCs, RGB, and so on, into the mainchain is not only probably technically impossible, but in a more conceptual sense – arguably an *aesthetic sense* – is just an obviously bad idea. It would introduce unknowable attack vectors and hence holistic fragility. The naïve view is that this compounds the utility of every functionality. The mature view is that it compounds only the vulnerabilities; each functionality is primarily affected to the extent it has become more vulnerable, and utility dramatically decreases, both at the level of individual functionalities and the protocol as a whole. If TCP/IP had been configured to enable video streaming, for example, it would have broken immediately if it had ever worked at all. This is a feature, not a bug. It reflects the mindset not of a cargo cult bureaucrat, but rather of a prudent and humble engineer, mindful of Gall’s Law from John Gall’s *Systemantics*, with which we opened this section.

As Gall’s Law suggests, we believe the general principle favoring Bitcoin’s layered architecture is not one of *software* engineering so much as engineering entirely in general, yet as elegantly applied to software. “*This clear specialization ensures performance, reliability, and scalability of the Internet,*” As Thibaud Maréchal puts it in [A Monetary Layer for the Internet](#). This might seem like an argument in favour of the likes of The Lightning Network from an oddly axiomatic basis – and almost a fatalistic one along the lines of: *software eats the money*. However, this rough idea has ample historical precedent that predates “software” by several centuries – probably precisely because the key insight is one of institutional design, transcending software entirely, and of which software is one special case amongst many.

One of the features of the complex web of financial and banking relations in Renaissance Florence was the practice of “offsetting”: noncash and *nonbank* payments between merchants by flow of credit and debit. Richard Goldthwaite describes in, [The Economy of Renaissance Florence](#), that,

“*One could draw on his credit by written order for transfer to a third party, and the transfer could be passed on to a fourth party and even on to others by mere book entry.*”

These “payment channels” were clearly private, and a final link to The Lightning Network is to realise this assumed a kind of *going concern*. In other words, that it was worth costlessly keeping credit channels open and updating them rather than closing them at cost, which would involve settling either in bank transfer, or with true “final settlement” in specie.

While the mechanical allusion is intriguing, Goldthwaite goes on to place offsetting amid the diversity of financial customs,

“*Local banks did not have a commanding position in the local credit market. On the supply side of the market, the weakness of these banks in attracting deposits was exposed by their failure to provide an outlet for the savings that began to accumulate in the hands of artisans and shopkeepers in the second half of the fifteenth century. The depositories opened by the Innocenti, Santa Maria Nuova, and the Badia, in contrast, responded to this void in the market, signalling the new direction banking was to take in the following century. But it is when we turn to the demand side of the market that we can see banks’ relative inability to attract capital. Local banks and especially pawnbrokers served the general public as sources for direct loans, but they were hardly the only conduit to credit. Direct loans were also readily available outside of banks. Evidence for loans from private persons abounds in the city’s oldest notarial records ... Moreover, debits and credits recorded in these*

official documents could be reassigned through another notarial act, although it is difficult to say that traffic of this kind constituted a secondary market.”

Although by no means Goldthwaite’s point, there is an obvious lesson from this historical analysis in comparing the merchant-driven, hard money economic system of Renaissance Florence to the finance-driven, soft money of modernity, and taken to the power of some fresh Hell by crypto. The lesson becomes even more helpful with an eye on a Bitcoin Standard near- to medium-term future: financial institutions and payment methods alike will mould themselves to the heterogeneity of time preferences, commercial requirements, and interpersonal customs to be found across society. There will not be “the bank,” gatekeeper to all finance, nor “the protocol,” carrier of all data and value. There will be a supply and demand of capital, liquid and illiquid, short-term and long-term, risk-seeking and risk-averse, financial and production, personal and professional, payment and settlement. Moreover, in Florence, this diversity of capital was priced and kept honest relative to the store of value of elemental gold. Gold itself was therefore disconnected from the likelihood of gradually debased coinage or even confusing alternatives for units of account. Gold was for final settlement, not for payment, credit, or capital. Of course, as effective and elegant as this system was, Bitcoin is *even better*.

In this light, Lightning, Liquid, DLCs, RGB, Stacks, and whatever else, are not clunky or bizarre in the slightest. They are natural, complimentary, healthy, and aesthetically and institutionally sound, as will be all other successful and differentiated extensions of the base layer. And in this light also, Curve, Aave, and so on, are deeply unsound: the entire Ethereum, Solana, Cardano, EOS, Tezos, Tron, etc., timechains carry and depend in part on the success of their crypto ecosystems. If Lightning is hacked and the liquidity drained, Bitcoin won’t care. The DAO hard fork, on the other hand, is a clear precedent, which, incidentally, happened to evidence violation of just about every one of Ostrom’s design principles for common pool resources.

In the *Odd Lots* episode mentioned above, Weisenthal made an astute observation just a little later in his introduction, saying,

“It’s not obvious to me that the people in the Bitcoin world want [yield farming and automated market makers]. It’s not obvious to me that people in the Bitcoin world want those things; whether there is an appetite or an interest for all that speculation and trading.”

He is *nearly right*. What is genius about Bitcoin’s architecture – arguably its *governance*, to foreshadow [Appendix A](#) for the interested reader – is that it doesn’t matter what people want or have an appetite for. Everybody competes for block space with fees and beyond that leaves each other alone. If RGB falls apart, Lightning users won’t care, and vice versa. If hardcore enthusiasts object in principle to Stacks stamping polluting the timechain, or to tokens or NFTs on Liquid, pegged to a DLC, both, neither, whatever ... then none of Blockstack, Blockstream, Suredbits, Lightning Labs, Liquid users, Lightning users, DLC users, anybody at all, nor Bitcoin itself will care. The purists are free to not engage with any of it and the reckless experimenters are free not to require the engagement of dissenters. Nobody needs to care because nobody is harmed by anybody else’s behaviour. Risk is sandboxed for those intending to take it. *Nothing is systemically important*. Imagine if regular finance had so magical a property...

All this loops back around to the Pfefferian argument first articulated in [Section 2](#), that there is reason to suspect that tokens in non-Bitcoin crypto that aim for similar functionality, *even if they work technically*, will have effectively unbounded velocity, negligible holding periods, hence no reason to sustain value, hence *actually* no reason to suspect they *will work technically*. The lack of an incentive to hold comes precisely from the tokens *not being money*.

Money is *universal credit*, not highly specific credit. Highly specific credit – to be used in a casino, or providing liquidity to an automated market maker – will lose in the market to universal money for painfully obvious reasons. Although the subject of a fascinating ongoing debate as to its applicability in light of Bitcoin’s emergence, this is a key element of the Mises Regression Theorem from, [The Theory of Money and Credit](#) – or rather, we can infer from the theorem an inverse non-implication: a “token” that is tradeable only for a highly

specific set of goods or services *will not become money* in the presence of less restricted but otherwise technologically comparable competition.

This argument can be graduated to a more general criticism that we can ultimately tie directly to Bitcoin. For most (and probably all) of the crypto projects in question, *there is no need for a token*. This might sound like a devastating critique but it actually contains the seeds of an olive branch. It is not necessarily that the ideas are bad in terms of the functionality or ethos aimed for; it is that the architecture supporting the current instantiation of the idea is poorly conceived and likely will not last. But it will likely be possible to rebuild most of it on Bitcoin. Layered architecture is the essence of the prospects for success of Lightning, Liquid, and so on. Or, in the framing just adopted: *they don't need a token either!* Which is great, because they don't have one. They use Bitcoin. Bitcoin is the universal credit on which their highly specific credit is built and to which it refers – or not-even-credit but assets that benefit from censorship-resistant, integrity-assured digital existence on a distributed ledger.

We stress once more that we do not directly endorse the technical or commercial merits of any of these projects and are even partial to the idea that the total failure of a few of them would be an excellent stress test for Bitcoin. Purely optically, that is. We know from first principles there would be no real stress, but for those less interested in first principles and more in hype and optics, this would be useful. When the Ethereum DAO was exploited, the entire community had to hard fork the timechain to reverse the catastrophic ecosystem risk. Bitcoin will never need to do this.

6/6 – Why We Might Be Wrong

“I don’t really care about being right, I just care about success. I don’t mind being wrong, and I’ll admit that I’m wrong a lot. It doesn’t really matter to me too much. What matters to me is that we do the right thing.”

Steve Jobs on “Crypto”

TLDR: *self-explanatory. Only the conclusion follows. We hope you enjoyed! :)*

This is clearly not an exhaustive list, but in the interests of intellectual honesty (and hence professional seriousness) we may turn out to be wrong for any or all of the following reasons:

- We have overestimated (or are simply wrong about) crypto’s technical flaws and underestimated its social strengths. Ever-growing inflows of capital, talent, and interlinked network effects may contribute to overcoming whatever flaws exist and the problems are successfully engineered around.
- Crypto programmability is never truly matched on Bitcoin and the value spiral and Pfefferian holding period problem alike are nipped in the bud by provably enforceable atomic swaps such that Ethereum, Solana, Cardano, EOS, Tezos, Tron, etc., effectively *become sidechains*, or ditto, but for obviously different reasons for Monero or ZCash, in which case the rest of crypto across either flavor of ecosystem can tack on too. wBTC arguably demonstrates pretty cleanly how this would work, except that what we suggest here would be *more robust still*.
- DeFi on Bitcoin just doesn’t work for whatever reason beyond the lack of programmability suggested above. Maybe all the projects above fail for technical reasons while their social utility has become clear and hence is serviced by crypto.
- Staking bootstraps financial viability on the basis of representing an *actual, provable* yield to ecosystem participants, short-circuiting what we described above as something of a *logical death spiral*, and forming a robust foundation for capital formation that we have argued is currently lacking.
- Bitcoin itself fails and everything we say here is relatively moot because crypto are all we have left. For the sake of argument (as it is hard to imagine Bitcoin fails but altcoins do not) let’s say this is due to some cryptographic vulnerability too grave to hard fork out and that crypto just doesn’t have. Or perhaps governments unanimously coordinate an all-out attack against Bitcoin that (as we know cannot happen) doesn’t destroy the network or disrupt its ability to facilitate savings and payments but *does* make it untenable as a basis for legal capital formation.
- Social factors: it all works but people don’t like using it because it’s too clunky. Again we feel this is unlikely but there is something of a precedent in the development of Web 2.0 – we could have all cared about privacy, sovereignty, and so on. Some did, and did so loudly, but it wasn’t enough because the vast majority preferred the slick utility of Facebook and Google.
- More of a caveat, but just to clarify we have only focused on crypto ecosystems that explicitly try to capture “decentralized finance” whether as competitive monetary base layers or as infrastructure for

capital formation. We do not focus at all on the likes of Filecoin, Sia, Storj, etc. that are derivative of “blockchain technology” but which aim to solve some completely different technical problem. Although it is certainly possible that many of our criticisms above apply to these protocols, they may succeed, for all we know, and their tokens may even attract a stable or growing value. We simply point out that we are not *intentionally* providing a critique of these assets.

[END] – Where We Go From Here

“It ain't what you don't know that gets you into trouble. It's what you know that just ain't so.”

Mark Twain on “Crypto”

Everything above is what we believe as professional capital allocators and technology enthusiasts. To borrow a wonderful quip from Alyse Killeen in the same *Odd Lots* podcast referenced a few times above, we are not Bitcoin maximalists, as it might seem natural to classify us in light of this piece; we are *sound tech maximalists*. We might cheekily go even further and say we are *intellectual honesty and responsibility maximalists* – which itself explains our fascination with Bitcoin.

And besides, there are plenty of technologies besides Bitcoin in which we are intellectually interested and towards which we will seriously consider advocating putting capital to work, contingent on all the usual factors of a sound business plan, a trustworthy management team, a realistic assessment of the technological and competitive landscape, and so on and so forth.

But for now, let us leave our thoughts at the following: we are excited by Bitcoin. But crypto? Not so much.

Thanks to Muneeb Ali, Andrew Bailey, Dhruv Bansal, Ben Carman, Tuur Demeester, Ryan Gentry, [Gigi](#), Joe Kelly, Alyse Killeen, Cory Klippsten, Theo Mogenet, Samson Mow, John Pfeffer, [Balaji Srinivasan](#), Elizabeth Stark, Alex Thorn, Robert Wilson, Reuben Youngblom, and Giacomo Zucco for edits and contributions.

It should go without saying that not everybody mentioned above agrees with any or all of our theses or conclusions. In fact, some passionately disagree, in which cases we appreciate their input even more greatly. Several contributors preferred to remain anonymous as well.

Appendix A - Common Pool Resources

TLDR: *in Appendix A, we argue that all “crypto assets,” Bitcoin included, are properly understood as “common pool resources,” as opposed to, for example, public goods. We then argue that, according to arguably the most respected analysis of such entities, the governance characteristics of Bitcoin are excellent while those typical of crypto are poor.*

An alternative justification for breaking the innovations in Bitcoin is often given along the following lines:

While crypto admittedly has some technical flaws, we should bite the bullet and look to its governance virtues instead. In this light, Bitcoin’s net flaws are even greater given Bitcoin effectively privatizes what ought to be a public good, leading to unacceptable governance.

In our view, this analysis is weak, and elucidates *another* avenue of analysis of the superiority of Bitcoin’s design, ecosystem, and community.

Is Bitcoin a public good? Are crypto tokens? *Is money?*

Bitcoiners may scoff at even the asking of the question and wonder if we have been drinking the modern monopoly money Kool-Aid. Surely money is private property? But the question is worth considering if only to answer in the firmly negative. As George Selgin quipped, in response to the [Bank of International Settlements seeming to think](#) the answer is “yes”:

“The argument that money is a “public good,” is one of many unfounded claims made about it that serve as “debate stoppers”: by uttering those magic words, experts hope to avoid having to otherwise defend state money monopolies.”

In her classic of political philosophy, *Governing The Commons*, Nobel (memorial) prize winner Elinor Ostrom gives a rigorous analysis of what she calls a *common pool resource* and the “problem” of governing its use. To be completely clear, we treat the following as an interesting analytical exercise and not a tool for slipping a line of thinking into our discussion that is subtly opposed to private property. Even the title of Ostrom’s book is potentially misleading in this regard; by “governing” she means something more like “decision making with respect to,” rather than “enforcing a decided-upon rule,” as the cognates of “govern” might unfortunately suggest. This would be a particularly inapt reading given her thesis – presented here so concisely as to absolutely not do it justice – is that there is a vast range of *common pool resource* [CPR] problems that are in theory, and have been in practice, better solved without government intervention and even without force of any kind but rather with effectively established communities, relations, and incentives. Also, often the same class of such problems have been made much worse with government intervention exercised by an authority entirely removed from what are essentially local issues. It is typical in such circumstances for the authority (*government* or otherwise) to arrogantly ignore, or perhaps even be entirely ignorant of, exactly such alternative and likely already existing methods for governing the commons.

In the book’s very last page, Ostrom laments what seems to her to be the default instinct of her academic colleagues in first, and often only, thinking of a government solution to any collective action problem. She writes,

“The models that social scientists tend to use for analyzing CPR problems have the perverse effect of supporting increased centralization of political authority. First, the individuals using CPRs are viewed as if they are capable of short-term maximization, but not of long-term reflection about joint strategies to improve joint outcomes. Second, these individuals are viewed as if they are in a trap and cannot get out without some

external authority imposing a solution. Third, the institutions that individuals may have established are ignored or rejected as inefficient, without examining how these institutions may help acquire information, reduce monitoring and enforcement costs, and equitably allocate appropriation right and provision duties. Fourth, the solutions presented for “the” government to impose are themselves based on models of idealized markets or idealized states.

We in the social sciences face as great a challenge in how to address the analysis of CPR problems as do the communities of people who struggle with ways to avoid CPR problems in their day-to-day lives.”

So with this critical clarification in mind, in what sense might money be a *common pool resource*? It is likely instructive to be clearer about how *exactly* Ostrom defines a common pool resource, and how she distinguishes it from a “public good,” as was previously canonically treated by Mancur Olson in *The Logic of Collective Action*. Ostrom writes,

“The relatively high costs of physically excluding joint appropriators from the resource or from improvements made to the resource system are similar to the high costs of excluding potential beneficiaries from public goods. This shared attribute is responsible for the ever present temptation to free-ride that exists in regard to both CPRs and public goods. There is as much temptation to avoid contributing to the provision of public security or weather forecasts. Theoretical propositions that are derived solely from the difficulty of exclusion are applicable to the provision of both CPRs and collective goods.

But one’s use of a weather forecast does not subtract from the availability of that forecast to others, just as one’s consumption of public security does not reduce the general level of security available in a community. “Crowding effects” and “overuse” problems are chronic in CPR situations but absent in regard to pure public goods. The subtractability of the resource units leads to the possibility of approaching the limit of the number of resource units produced by a CPR. When the CPR is a man-made structure, such as a bridge, approaching the limit of crossing units will lead to congestion. When the CPR is a biological resource, such as a fishery or a forest, approaching the limit of resource units not only may produce short-run crowding effects but also may destroy the capability of the resource itself to continue producing resource units. Even a physical resource, such as a bridge, can be destroyed by heavier use than was allowed for in its engineering specifications.”

This potentially gets dangerous again in terms of the blurring of what is and is not truly and unmistakably private property, hence the all-important caveat cited above. But we think that a more pragmatic analysis of “money” forces us to move beyond what we might call its “ideal qualities” and realize that, in real life, money has always been both private property *and* affected adversarially by the “subtractive” behaviour of others, free-riding in a manner that *clearly* harms the well-behaved. As will hopefully clear up any doubts as to the clout of this line of argument, consider Mises’ line in [his pithy tract on “velocity”](#) that, “*money is of course a social institution.*” Private property is not: it is arguably an *anti-social* institution. It exists in a state of nature so plainly as to make demarcation moot and is only potentially infringed upon *by society*, and the need for interpersonal compromise that society introduces to human affairs. “Money in a state of nature,” on the other hand, is entirely nonsensical.

Lawrence White helpfully demurred on the above to the effect that surely money *balances* are not a common pool resource given Alice cannot spend Bob’s balance and vice versa, except without undeniable appropriation. We agree, and, in fact, think this clarification bolsters our own claim. Money balances are private property but the *institution of money* is a common pool resource, in arguably *exactly* the same way that a stock of fish might be a common pool resource, even though fish *fished* by fishermen have clearly become private property.

We see no reason to take “depletion” only literally. We do not mean that physical coins or notes depreciate, which would ironically have the opposite effect on the value of the common pool resource in this case. We mean that the *utility* of the institution depletes with inflation. And why do we care about stocks of fish and their possible depletion if not for the economic utility of fish? We are not pescaphiles. Within the confines

of this discussion, at least, we do not care about the spiritual value and beauty of fish in the wild. We care about the depletion of fish stocks due to the economic calculus of acting and interacting humans, not the random vandalism of marauding pescaphobes. Were there not an incentive in the first place to deplete fish stocks by fishing, there would be no need to classify it as a common pool resource in need of effective governance.

What is seigniorage if not the depletion of a common pool resource by a nefarious “crowder”? What is gold-mining if not a less nefarious and admittedly costly activity (so not exactly “free riding,” either), but nonetheless a depletive interaction with a common pool resource? And *what is this common pool resource* if not something ultimately psychological and reliant on subjective value? Is this not precisely Mises’ point above? Money is not the coins, or the balances, or even the UTXOs, but the consensus around economic behaviour and reality these “tokens” capture. We think it is neither a philosophical stretch nor an authoritarian backdoor to say that, yes, money as an institution *literally is* a common pool resource, and that *what money is or should be* is therefore a common pool resource problem. It may be the single most important common pool resource problem; hence Bitcoin is the most important solution, and by extension the most important enabling technology for the management of a common pool resource.

This obviously distinguishes money from, for example, a far more prosaic private good like a mug. Alice using Bob’s mug clearly prevents Bob from doing the same. The good is rivalrous. Bob can prevent Alice from using his mug by stealing it, breaking it, etc. (in effect, incurring a *tort*). However, money occupies a nebulous middle ground between cleanly rivalrous and non-rivalrous: if Alice prints her own money, Bob is likely none the wiser as Alice hasn’t *stolen his coins*, and there is clearly no comparison to be made to Alice *printing Bob’s mug*. And yet Bob’s money has been adversarially and extractively affected, because it is not the token that matters, it is the consensus the tokens represent. Which is, of course, to say that money as an institution is not a private good, but a common pool resource.

Great. So what?

So, in *Governing the Commons*, Ostrom identified eight “design principles” for analysing the governance merits of CPRs. Below, we go through them one by one, occasionally offering some commentary on the applicability of the principle in this space, and evaluating in the case of Bitcoin and altcoins. Instead of treating altcoins as homogeneous, given they come in a variety of forms, we thought that looking at Ethereum would make the most sense. The reader is encouraged to keep in mind, however, that likely *most* of what we say about Ethereum likely applies to all of crypto-ex-Bitcoin – just not *all*.

I - Individuals who have rights to withdraw resource units from the CPR must be clearly defined, as must the boundaries of the CPR itself

*There is a latent ambiguity here given Ostrom didn’t anticipate a CPR that allows for differentiated modes of “usage” on the one hand and “depletion” on the other: she assumed these meant the same thing. In our case, “use” means what it sounds like, whereas “depletion” means token inflation. We cover both in this point, answering: **are rights clearly defined?** – but note in later points only “token inflation” is really intended by her subsequently introduced terminology of “appropriators.”*

Use

Bitcoin: yes, invalidly signed transactions guaranteed to be rejected (integrity assurance), and validly signed transactions are overwhelmingly likely to be accepted (censorship resistance) unless miners *attempt* to censor, which is overwhelmingly likely to fail in the long run.

Ethereum: largely, yes. But not “unequivocally, yes” because there are constant updates to the protocol that users are powerless to reject, as well as an adverse precedent in the hard fork following the DAO hack that, if repeated, would further call this point into question.

Inflation

Bitcoin: yes, block subsidy schedule is transparent and highly unlikely to ever be altered.

Ethereum: largely, yes, but again some caveats. There once again are changes that are currently underway that will change the inflation schedule as well as criticisms that current supply is unreasonably difficult to verify.

II – Congruence between appropriation and provision rules and local conditions

Ostrom's original point here was that we must always adopt a localized and hence flexible mindset and not design CPRs from an armchair. It's a bit of a non sequitur in this case as crypto-assets are inherently global, but we could perhaps consider the "local environment" not to mean anything geographic but rather relating to something like an "Internet ecology."

Bitcoin: irrelevant as money applies to everything so there are no "local conditions," no matter how "local" is defined.

Ethereum: not quite as strong as Bitcoin but likely similarly irrelevant given "computation" is general enough to not impart any importance to "local conditions."

III – Most individuals affected by the operational rules can participate in modifying the operational rules

Bitcoin: yes, near consensus of users is required for rule changes.

Ethereum: not strong enough, in our opinion. We see clear authority wielded by a few core investors and developers. Also we have concerns that as Ethereum moves to proof-of-stake this concern could become even worse as the "oligarchs" of the protocol will retain even more control – bizarrely, entirely by design.

IV – Monitors, who actively audit CPR conditions and appropriator behavior, are accountable to the appropriators or are the appropriators

Bitcoin: yes, and probably the single best designed CPR in history in this respect. Anybody can moderate "appropriator" (i.e., inflator, i.e., miner) behavior for effectively no cost, and this auditing makes it effectively impossible for appropriators to do anything at all other than what is allowed, which incidentally *benefits* the CPR by providing security and so is questionably even "appropriation," at least not in precisely the sense Ostrom intended.

Ethereum: yes and no. Yes in theory, given Ethereum runs on a similar proof-of-work algorithm (for now). But no in practice given the data structure is so large and the bandwidth required to keep up to date so high that very few monitors even exist, and a large proportion of those who do rely on third party infrastructure providers to whom they are arguably beholden, hence not independent and accountable to parties in addition to the appropriators.

V – Appropriators who violate operational rules are likely to be assessed graduated sanctions (depending on the seriousness of the offense) by other appropriators, by officials accountable to these appropriators, or both

Bitcoin: yes, and probably the single best designed CPR in history in this respect. Appropriators cannot appropriate outside the rules and, if they try to, other rule-abiding appropriators will instantaneously benefit at their expense.

Ethereum: yes, very similar to Bitcoin, with the minor caveat that a far lower proportion of appropriators are required to change the rules (i.e. in violation of what everybody else *thought the rules were*) than in Bitcoin.

VI – Appropriators and their officials have rapid access to low-cost local arenas to resolve conflicts among appropriators or between appropriators and officials

Bitcoin: yes, and probably the single best designed CPR in history in this respect. “Officials” are users, and the conflict resolution mechanism is thermodynamics and math.

Ethereum: we see some issues here where conflict resolution has seen to come down to a handful of overly powerful individuals. And again, we see proof-of-stake aggravating these issues.

VII – The rights of appropriators to devise their own institutions are not challenged by external governmental authorities

Bitcoin: yes, and probably the single best designed CPR in history in this respect. Governments could in theory interfere with Bitcoin, but the possibility is diminishingly likely and becoming more so every day. As opposed to the potential to challenge any other CPR in the world, or in history, the cost to “challenge” Bitcoin is so much higher as to be basically unfathomable and incalculable. Ostrom’s *intended* point likely doesn’t even apply – or does apply but becomes something of a technicality.

Ethereum: superficially, yes, but highly questionable in the long run, for two reasons. **Most** of the full nodes are run in centralized data centers (easy targets for a government – or, for that matter, any adversarial actor – we go into this in more detail in [Section 3](#)) and individuals who hold both disproportionately large voting share and community-weighted influence should they be externally compromised.

VIII – Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organized in multiple layers of nested enterprises

Bitcoin: yes and no, depending on how to interpret the design principle. No, in the sense that there is only one function, one way to conceive of attempted illegitimate appropriation, and one conflict resolution mechanism, which is likely not even a bad thing in terms of this design principle, but simply highly unusual as CPRs go. But yes, in the sense that Bitcoin functions as a base layer for a range of layered CPRs that are rooted in its functionality, which can indeed be thought of as “multiple layers of nested enterprises” and, it could certainly be argued, are a crucial part of Bitcoin’s functionality. We discuss this in much more detail in [Section 5](#).

Ethereum: yes and no, for essentially the same reasons as Bitcoin. The discussion we allude to above in [Section 5](#) is more about technical functionality than governance.

Obviously, we do not claim that Ostrom’s framework is the absolute truth of this matter. Readers are entirely free to ignore her design principles or our analysis and application of them. Our intention is not to suggest the unquestionable validity of either, but rather to point out that hers is the most respected analysis of *the kind of thing* Bitcoin and altcoins are, and that it is certainly useful at least comparatively.

If the claim is to be made that altcoins may have various technical inferiorities but that they make up for these with superior social characteristics on the basis of better or fairer (or whatever) governance, we would strongly disagree. Not only are they not better, they are markedly *worse*. They invite attack. They practically demand it at the precise moment the governance mechanisms start to creak.

Appendix B - Rehypothecation Algebra

TLDR: *fuller explanation of the algebra cited without workings in [Section 2](#).*


The example given in [Section 2](#) was 50% overcollateralization, or 150% collateral. This means we require \$150 as collateral to mint \$100 of new assets. However, if these can be reused as collateral for another such operation, we apply the same ratio $(100/150)=2/3$ to 100, to see that we can mint another \$67, and so on, ad infinitum. It takes two iterations of this for the total $(\$100 + \$67 = \$167)$ to be higher than the original collateral, and the limit of outstanding asset value approached can be calculated as follows:

$$\$100 *_{n=0\infty} (2/3)^n = \$100 * 3 = \$300$$

In general, then, if the collateralization ratio is $x\%$, for $x \geq 100$, then let $k = (100\% / x\%)$, and our formula becomes,

$$\$100 *_{n=0\infty} (1/k)^n = \$100 * (k/k-1) \quad \text{for } k < 1$$

The boundary condition on the right captures that a collateralization ratio of exactly 100% or lower clearly leads to an infinite sum (this is easier to understand practically than via the equation) and hence the equation will pop out either a practically meaningless (possibly negative) answer, as can be seen by charting $f(y)=xx-1$ below:



It is easy to see that the single point at which $f(y)=y$ is where $y=2$, or the capitalization ratio is 200%. This was alluded to in the main body of the paper and is also easy enough to understand practically, as the reuse of the products of iterative 200% collateralization would look like putting up \$200 to generate $\$100 + \$50 + \$25 + \$12.50 + \dots$ easily recognizable as converging to \$200.

Collateralization ratios above 200% are even “safer” in the sense of unboundedly many reuses of the proceeds will never reach a synthetic exposure that exceeds the initial collateral, or, in more technical terms, $f(y) < y$.

It is where $f(y) > y$ – so $y < 2$, but also where the equation is inside its boundary conditions, hence, $y > 1$, where things get interesting. These are collateralization ratios where the total value of asset issuance following reuse of the product of collateral will eventually reach a higher value than that of the initial collateral. We highlighted in the main body of the paper that, for example, 175% takes 3 iterations to surpass the initial collateral and approaches a limit of \$233.