The Bitcoin Block Size Limit, Artificial Scarcity, and Code-Enhanced Public Club Governance

Konrad S. Graf

24 December 2019[1]

1. Introduction

How can Bitcoin's block size limit and associated decision-making processes be characterized in economic, legal, and social-theory terms? In mid-2016, I advanced several observations on what I called Bitcoin block size political economy. The current paper covers developments in my analysis since then, revisiting these topics to place them into several broader perspectives. It also follows up on my review essay on Ammous 2018 (Graf 1 Oct 2019).

To the narrow extent that a "market test" can be considered to have happened with regard to the block size limit—narrow because so many other variables are also present—the Bitcoin (BTC) approach has succeeded over competing approaches, as exemplified by the price performance of the Bitcoin Cash (BCH) chain split. BCH currently trades at under 3% of BTC, with a similar relative proportion of hash rate. This stark outcome, though it follows from a constellation of many complex factors, not only one, nevertheless deserves examination, to which the several concepts and models below should also contribute.

I continue to aim for as descriptive an approach as possible in examining how Bitcoin's qualities and dynamics can best be characterized, particularly as interpreted from economic and legal viewpoints grounded in the action-theory approach of the Austrian school. Key concepts will include: 1) the differentiation of the *transaction-inclusion market* from the non-market for *verification & relay services*, 2) *voluntary-sector artificial scarcity* versus both natural scarcity and compulsory-sector artificial scarcity, 3) *code-enhanced public club governance*, or the competitive promulgation of rule sets for non-state institutions as mediated by technology for defining, maintaining, and enforcing them over time, and 4) the application of evolutionary models to ideological orientations, foremost, in this case, hard-fork avoidance versus hard-fork embrace.

2. Review of previous observations

Among my earlier observations in an interview (Graf 4 May 2016) and series of follow-up posts (8, 9, and 10 July 2016), I argued that once the block size limit began to function as an active restriction on normal transaction volume for the first time, economic and sociological dynamics new to the system would ensue with both immediate and longer-term consequences easy to miss or underestimate. Up to around that time, average transaction volume had remained below the limit, leaving it for the most part economically inconsequential, much as a minimum wage well below the lowest level of market wages typically paid has little practical effect.

Most immediately, I argued, the limit would begin to operate as a production ceiling or quota on a services market, with certain effects on pricing and volume. Structurally, it would promote the growth of off-chain services at the expense of on-chain transacting, altering the course of competitive industry evolution, and, in effect, subsidizing the development of off-chain alternatives. Sociologically, it would become a "political" object of controversy, with a divide emerging around differences in values and priorities for the system as well as differential financial incentives stemming from the positions of various business models vis-à-vis the production ceiling's effects.

A common argument in favor of the limit was and is that it would boost mining revenue from transaction fees, enabling this revenue stream to take over from the quadrennially halving fixed block reward ("miner subsidy") sooner than otherwise. But given the several drawbacks of such a limit, could other methods enhance miner funding instead? To consider system financing under a higher limit, I looked to the system's physical and energetic points of natural scarcity. I considered mining business decision-making (*transaction-inclusion services* market) and possible methods of direct node financing (the introduction of pricing for *verification & relay services*). I wondered if these could not come to constitute sufficient bases for supply & demand dynamics without need for an output ceiling maintained at an apparently arbitrary level.

I had also argued earlier that off-chain transacting media such as sidechain coins and payment-channel network units could not "be" bitcoin, as some promoters claimed, but instead constitute *bitcoin substitutes* (Graf 24 Oct 2014). Whether these units could come to function as *perfect substitutes* could only be seen in practice. In essence, to reach this high standard, a typical market actor would have to be indifferent as to whether they were paid in on-chain bitcoin or with a given substitute. A market outcome of broad acceptance of substitutes beyond a sphere of hobbyists and enthusiasts could not be assumed based on technical intentions. A technical peg would not automatically result in open-market parity. That would only be one possible result and a sign of a successful *outcome*, one that could not just be assumed in advance.

The combination of the foregoing led me to some skepticism at what appeared to be a movement to entrust Bitcoin's future to products under development that *might* work, in contrast to on-chain Bitcoin, which *did*. Raising the limit appeared to be a practical way to maintain progress while more experimental systems could come online and be market tested in due time. I now follow up on these observations with additional models that place them into wider contexts.

3. Two apparent tragedies of the commons

The immutability of data recorded on the block chain and decentralized trustless verification of additions to it are part of the core of Bitcoin's value. However, rising data throughput increases the costs of running full network nodes engaged in verification and relay of new transactions and blocks. Factors include the accumulating size of the chain, the varying size of the *unspent transaction output* (UTXO) set, the size of the *mempool* (the collection of outstanding transactions bidding against each other for confirmation priority), the size of new blocks, and resulting overall bandwidth demands.

However, methods by which full node operators can earn *directly* in exchange for services are absent, leaving node operators less able than they might otherwise be to respond with price signals to the burden of higher throughput, let alone chain-size growth. Node operators are left with a far blunter primary instrument than pricing—to either continue running or shut down in the face of a partial commons of traffic.

The *transaction-inclusion market* as identified in my earlier analyses concerns relationships between miners and transaction senders, but as I also then noted, does not directly address or coordinate feedback between traffic volume and chain size and the provision of full-node services. Full nodes act as relays that bring together transaction-inclusion bidders with suppliers (miners). Despite being in this seemingly opportune position, however, node operators face practical and technical limitations in monetizing their role through direct fee collection. We might refer to this as the *verification & relay market,* except that Bitcoin's design supports no such market. These market demands are met either through volunteer contributions or as a byproduct or operating cost of other activities.

Some other cryptocurrency designs, such as Ethereum and Dash, had certain direct node-funding methods built into their frameworks. However, it may or may not be possible to build direct node funding into Bitcoin—or judged desirable on balance to do so. In addition to technical and structural challenges of selecting and introducing any such methods, direct compensation could also lead toward the automatic labelling of nodes as commercial entities, potentially exposing them to more burdensome tax and regulatory classifications than they could enjoy as volunteers. Such a result could tend to depress node count, the opposite of the sought-after effect.

Instead, Bitcoin node count and quality have relied on—and may have to continue to rely on—a blend of ideological and altruistic voluntarism and non-altruistic cross-interests. For example, running a full node is an operating cost of providing other Bitcoin services, from payments to trading to data analysis to mining. Running a full node also brings certain direct benefits to the operator, such as unmediated verification, increased address privacy, and lower-latency network connectivity. All of this still leaves no *direct* node services compensation, only voluntary and indirect motivations. If one counterfactually imagines that some direct pricing model could better coordinate demand for with supply of verification & relay services, this absence can be viewed as setting up a partial *tragedy of the commons* situation.

More fundamentally, though, the block chain itself is a non-scarce good, meaning that it consists of pure information, freely copiable. As such, it does not naturally invite economic limits to its own expansion. Nevertheless, even though the chain—as pure

information—is a non-scarce good, the particular rate of its progressive expansion differentially burdens services and processes that are scarce: storage, bandwidth, and processing power. Chain size growth might thus be characterized as a second distinct commons problem, even beyond the more immediate issue of traffic flow.

The flow of transaction traffic might be met with transaction-fee pricing also in the absence of a block size limit, but the incentives behind it only address miner business decisions *on* transaction-inclusion—the business impact of the size of their own candidate blocks relative to transaction-fee value contained. This calculus does not fully extend to node-operator decisions about verification & relay services. Moreover, *even if* verification & relay pricing could be introduced, which seems doubtful, this would mainly influence the flow of traffic, the first commons problem, while not further and in addition addressing total chain size, the second commons problem.

Traditional property rights solutions inapplicable

In traditional analyses of commons situations, a recommended resolution is the development, definition, and protection of property rights in the resources in question, which leads to improved long-term stewardship and waste-reduction. With clear property rights, specific parties become beneficiaries of the discounted expected future value of specified resources, not merely the expected value of their immediate exploitation. Unowned resources are easily subject to a competition of which party can consume or otherwise exploit them first, rewarding immediate consumption over long-term preservation and sustainability. In Bitcoin's case, however, since blocks, transactions, and units are all pure information, the ownership model cannot be imported as a valid solution. Doing so would be a category error with serious repercussions (Graf 2015, 3–4). Other models are required to characterize and understand how Bitcoin is addressing these challenges.

Much of the order found in markets results from private, or non-state, governance. Stringham (2015) argues that the rules of even the most complex institutions and practices, including stock and later derivatives markets, emerged through a competitive process of trial & error rule selection within privately operated clubs and institutions. Rule sets, like qualities of a product, became part of an active competitive landscape. Some rule sets would enable certain institutions to thrive more than others, something discovered in specific implementations for definite purposes. Total cost/benefit balances of different rules would be unclear *ex ante*, requiring iterative experimentation.

Not only were such rules not originally imposed by benevolent governments supposedly looking to create order so that otherwise anarchistic private markets could then function better, such governments instead opposed and obstructed new rule sets step by step until finally capitulating only many years later to some of the valuable orderings that private institutions had already been succeeding with. Eventually though, governments went beyond acquiescence to begin forcing rules that some private institutions had originally invented, on other private institutions for which they were inappropriate. For instance, the specific rules for a tier-one stock exchange must differ in some respects from those for a higher-risk start-up market. Forcing teir-1 rules on a high-risk start-up market does not make it safer, it just kills it.

Governments, in contrast to non-state institutions competing along a spectrum of niches and applications, tend to foist inefficient one-size-fits-all rules with inferior regard for context and comprehensive cost considerations. This is natural because the costs that follow from governmental decision-making, including democratic voting, fall overwhelmingly onto parties other than the decisionmakers themselves. In contrast, non-state promulgators of rule sets tend to face the results of their chosen rules far more directly and comprehensively as organization members.

Stringham's fifth chapter discusses the evolution of rules in exclusive clubs in particular. Such institutions as trading coffee houses and later the London Stock Exchange, came to internalize both the costs and benefits of their rule choices, membership codes, and enforcement processes to participants, enabling an ongoing dynamic through which rule sets could be subjected to practical use tests in different applications.

Whereas property rights internalize the net present and discounted future value of ownable resources to owners, private rules can internalize the net present and discounted future net value of unownable rules to organizational participants, including the total package of real costs and benefits. Key innovation steps in the earliest stock markets included rules for exclusion. Prevention of default, for example, is far superior to addressing defaults after the fact. This was reflected in rules and judgment processes concerning who could and could not participate in a market to begin with, as well as procedures for banning defaulters. Early state efforts to obstruct private rule-making in nascent stock markets included attempting to force private trading clubs to accept members they had wished to exclude as actual or potential defaulters.

Bitcoin's design has taken this spirit of problem-prevention through *ex ante* rules even further. "There's no reliance on recourse. It's all prevention (Nakamoto 15 Nov 2008)." The code makes stringent demands of all participants in the form of its consensus rules. But in contrast to the traditional model of exclusive private trading clubs, whoever follows Bitcoin's rules can participate without the need for any human judgment to include or exclude particular members. Those who follow the consensus rules are in, those who do not are out. Participant inclusion/exclusion is moderated deterministically via consensus rules reflected in software. We will examine these concepts in greater depth in Section 5.

4. Two distinct cases of voluntary-sector artificial scarcity

The two partial commons issues defined above combined with the difficulties of introducing direct node funding methods to help alleviate them, may offer a reasonable basis for developers and other participants to support other measures to restrict traffic and chain growth, even seemingly arbitrary limits that enforce *artificial scarcity* of "block space" for practical ends.

The creation of artificial scarcity with the aid of law, to which we are accustomed today through intellectual property (IP) legislation, can be characterized as *compulsory-sector artificial scarcity*. The police power of the state enforces the creation of scarcity where it would not naturally prevail, for the profit of some at the expense of others.

This should be opposed on principle as accomplishing the opposite of what a legitimate system of property rights should, which is addressing the problems of *natural* scarcity

in a social context (Graf 2015, 55–60). Natural scarcity, and more specifically, *rivalness*, is given by the fact that two parties cannot make use of the same resource at the same time without coming into physical conflict in the attempt. In addressing such unavoidable issues of rivalness, various principles and rules of thumb can be employed. Enduring property rights based on objectively identifiable connections between owners and resources are a highly effective means of preventing and alleviating conflict and enabling durable foundations for social cooperation and economic prosperity (Hoppe [1989] 2010).

Compulsory-sector artificial scarcity, in contrast, begins with goods that are not naturally rival and attempts to make them so by law. Where more than one party could use the same good at the same time without coming into physical conflict, IP law declares that they cannot. In creating avoidable issues through artificial scarcity, IP law exacerbates conflict.

In the context of an all-volunteer system such as Bitcoin, however, this legal-theory objection to artificial scarcity measures does not apply, much as a property-rights solution does not apply to unownable goods. Indeed, a central contribution of Bitcoin is the novel way that it sets up artificial scarcity of units (monetary policy) such that no party is in a position to alter their production schedule. But could another, quite distinct, example of artificial scarcity, that of block space, also be advocated on other grounds?

Market processes display a natural tendency to route around efforts to create artificial scarcity, meeting it with abundance, the opposite of what anti-market ideology would have us believe. When pitted against compulsory-sector artificial scarcity, this routing around takes the form of ignoring, circumventing, or even contractually bypassing some of the effects of IP legislation, such as through open-source or Creative Commons licensing. When pitted against voluntary-sector versions, it takes the form of offering competing substitute products, subject to prohibition of fraud.

No one can sell knockoffs under the same brand name without risking a legitimate legal response on the basis of fraud. It is fraud if customers are led to believe they are consenting to buy one thing when they are in fact to receive a different thing. However, functionally quite similar substitutes can still be sold under other brands without fraud. Competitive production of close substitutes, barring fraudulent representation, pushes back against the natural quest for limited-edition pricing premia within a voluntary-sector context. Bitcoin has demonstrated striking resilience against market processes that degrade artificial-scarcity premia. Even in the absence of legal favoritism, it is able to maintain a spectacularly high "limited edition" premium for its units in a sea of thousands of knockoffs and variations.

This is due in part to peculiarities of its monetary context, in which value tends to gravitate toward a dominant asset based on network effects, market expectations, and—more concretely—Bitcoin's superiority in network hash rate and other security characteristics. A unit of any other cryptocurrency on *its* chain is a poor substitute for a unit of bitcoin on *its*. Although expressed as pure information, each bitcoin unit is inextricably tied to the specific use of naturally scarce economic goods—processing power, energy, and bandwidth—as applied to its chain in particular.

The obvious case of voluntary-sector artificial scarcity in Bitcoin is the strict scarcity of its transferrable monetary units. However, the block size limit constitutes another such case of artificial scarcity, one that is of a quite different character. Since any number of monetary units, provided sufficiently divisible, can produce the same services for a society of money users, increasing the production rate of such units adds no net value, and even subtracts value due to net losses from *Cantillon effects* and uncertainty about the unit's future inflation rate. In contrast, "block space" represents the total amount of on-chain data throughput per time period, which directly influences the quantity of a particular productive service that can be provided to paying customers—that is, on-chain transaction inclusion. Economic theory can fully support fixity in the quantity of monetary units, but fixity in the total quantity of a particular service provided across an industry of competing suppliers enjoys no such unreserved initial endorsement.

Nevertheless, since the block size limit is an instance of voluntary-sector, not compulsory sector, artificial scarcity, maintaining or altering its height is unobjectionable in legal-theory terms and therefore subject only to the forward-looking entrepreneurial and technical judgments of relevant participants. Such participants have a natural interest in designing or running software that has reliable restrictions—rather than merely partial or hypothetical ones—on the growth of commons resources that they will have to work with in the course of their opt-in participation, including traffic flow rates and accumulating chain size.

Previously, I had examined the fringes of natural scarcity in the Bitcoin system as potential sources of enhanced market mechanisms to constrain average block sizes through pricing rather than a fixed limit (Section 2). My key shift now is the assessment that Bitcoin's design is fundamentally reliant on multiple examples of voluntary-sector artificial scarcity, not only its headline example, the coin-production schedule. Since the edges of natural scarcity impacting average block sizes are merely peripheral, certain rules that support this with additional artificial scarcity are key to ongoing network success. This shifts focus to the content of such rules and the levels at which they are set, a judgment call most naturally led by those engaged with production practice, such as developers, technical experts, and hardware operators.

The use of artificial scarcity for maintaining a monetary policy is unproblematic for reasons unique to money. This use may therefore not give rise to the same challenges and controversy dynamics as the contrasting block size limit example has. Nevertheless, commons problems regarding use of verification & relay services and chain size, which the transaction-inclusion market addresses only peripherally, may necessitate ongoing reliance on voluntary-sector artificial measures. What kind of decision-making processes influence and shape such measures in the unique context of Bitcoin?

5. Action, consensus, and voluntary-sector regulation

The Bitcoin network's active code changes or remains in line with an emergent human consensus of relevant participants, including developers (which code to release), miners (which chain to mine on), full node operators (which chain to operate on), and exchanges (which coins to list). It does not follow any business- or government-like formal decision-making process. Critically, it is not even really a "consensus" in the common sense of "let us all sit down and agree to agree." At bottom, it is more action-oriented than argumentational or promissory. It boils down to developers of open-

source free software specifying the content of their releases, users selecting which such free software to run, if any, exchanges and wallet providers deciding which coins to list or support, and the net systemic results of all such discrete voluntary acts. It constitutes an emergent order. Argumentation and rhetoric can influence consensus-related actions, but the consensus arises from the actions themselves rather than any perceived or recorded discussion outcomes. This model also encompasses exit in that the human Bitcoin consensus reflects consensus among remainers.

The term *consensus* can also be confusing because *consensus rules* refer to those elements within cryptocurrency software that define constraints for what will be accepted as valid pieces of data, including blocks and transactions, on a given chain. A computer-science meaning of *consensus rules* thus co-exists with a human sense of *consensus* that is itself a peculiar variant of the ordinary concept.

None of these various voluntary acts of software use and selection, nor any of the various senses of *consensus* described above, constitute markets. As such, the expression, "the market's choice," can be misleading when applied to free software, such as when it is alleged that "the market" has chosen this or that implementation. It is a common error, shared by both critics of "the market" and at times some of its proponents, to divide the world into the compulsory sector of the state and the market sector of buying and selling. This tends to reduce the non-compulsory, voluntary sphere to a narrow view of "rational" transactional relations.

However, buying and selling comprise but a specialized part of the vast scope of the voluntary sector. Even among business entities, buying and selling operate within a far richer context of communication, understanding, tradition, long-term relationships, inertia, culture, private governance, and innovation. Rather than it being natural that pricing can directly mediate every voluntary-sector interaction, it is more the case that out of the vast scope of the voluntary sector, pricing can usefully mediate a few specific types of interactions.

Prior to the advent of pricing and markets, quite different methods for addressing tragedy-of-the-commons and other problems operated. These often took the form of tribally based management, traditions, and mores related to common-resource use. Pricing and markets are specific developments in the way humans have come to address certain economic problems. They are not the only types of voluntary-sector solutions, and for many applications, remain inappropriate.

In light of the foregoing, when faced with the two partial tragedy-of-the-commons gaps previously described—key spaces where direct market pricing is absent and may be impractical or impossible to establish—it would make sense if one found that the consensus, both human and encoded, reflected some other non-price method of limiting throughput to suppress the under-checked rise of verification costs and chain size. And what might that be?

The 1MB block size limit, already in place since 2010, has come to function as a fallback in this role. Leaving it in place as average transaction volume growth eventually climbed up to be impeded by it has left a rather blunt and imperfect check, but still one effective check, on the secular rise of verification costs and chain size. It may be understood as a

way of addressing by other means a problem that would be difficult or impossible to solve with a pricing-mediated supply & demand dynamic.

The on-chain transaction-inclusion market, with this limit on it, is able to ration artificially scarce block space among users through competitive bidding. This market is, in effect, "regulated" with an industrywide production ceiling applied to miners. Nevertheless, given certain difficult-to-resolve commons dynamics, as well as other perhaps even more significant factors to be discussed in Section 7, this may be on net the best strategy actually available, even in light of its several drawbacks.

An alternative strategy, for example, might be to set a minimum price for transactions, but this also raises intractable problems of arbitrariness and change over time as to the specific height of the minimum price. Compared to this, the existing volume-based limit has the easy-to-underestimate advantage of already existing. It also offers greater engineering-level predictability to node operators, miners, and software developers alike as to the future course of data flow requirements and chain size. This strategy is not without negative aspects, but no strategy is. Particular sets of costs and benefits are at stake with every possible strategy.

What is important from a systemic perspective is that the purveyors of rule sets, as well as those who chose to be subject to them, are the ones who face the brunt of both the costs and the benefits that follow (Section 3). This contrasts with state-style governance in which the selection of rules and the incidence of their costs and benefits are systematically separated, leaving rule selectors free to impose the costs of their preferred rules *onto others* who have neither selected those rules nor chosen to be bound by them. Bitcoin participants themselves elect to be bound by its current consensus rules, modulating their involvement via role types such as end-user, node operator, miner, developer, or peripheral business operator. In this choice, participants accept a rule-set package that includes a balance of tied costs and benefits.

The evolution of formal markets can bring vast increases in efficiency and extension of the scope of cooperation among strangers, but this does not happen through a "pure" and disembodied introduction of supply & demand dynamics everywhere one looks, like the proverbial man with a hammer seeing only nails. Instead, it is mediated by social expectations, mores, conventions, and systems of private governance.

Bitcoin encodes a set of club rules that apply to all who wish to participate—*because* they wish to participate. Participants benefit from certain *club goods* (Buchanan 1965). For example, use of the chain as a whole is partly non-rivalrous in being shared by all users as a common resource and partly excludable in that consensus rules place general limiting conditions on chain additions.

Bitcoin participants constitute more of a *public club* than a private one in the sense that Bitcoin's rules form the conditions on which anyone anywhere can participate or abstain. Rule enforcement and grounds for participant acceptance are already baked into the system itself and require no additional social measures or locality of relationship. Violation of consensus rules likewise automatically results in exclusion without any special human intervention. Yet while public in the sense that anyone can freely join without having to apply, it is in no way run by a state, making it part of a public sphere independent of the state.

Bitcoin brings a set of voluntary-sector rules beyond the context of a local tribe, private club, or particular market organization to compete on a global stage. But unlike the state, Bitcoin does not seek to apply its rules to any who do not opt in to be subject to them. Other similar rule sets (other cryptocurrencies) freely vie with Bitcoin for attention and membership, albeit with limited success.

Critically, Bitcoin's voluntary-sector artificial scarcity measures do not merely hang in an abstract, wishful realm, but are enforced via employments of naturally scarce capital equipment and energy. It is voluntary-sector artificial scarcity inextricably wrapped in a protective layer of natural scarcity such that ordinary market processes do not degrade the artificially scarce values in the ways they ordinarily would. The market value of the (artificially scarce) coin and the strength of its (naturally scarce) energetic-computational shell vary together.

Nevertheless, the shell of natural scarcity alone is insufficient to shape and maintain the specific structures of artificial scarcity being protected, for which various non-market but voluntary sector measures remain necessary to the valued package. Bitcoin's is a private governance model enhanced with rules embedded directly into the running system, *code-enhanced governance*. Combining these terms, we describe Bitcoin has having a system of *code-enhanced public club governance*.

6. Bitcoin as contrasted with bitcoin substitutes

Central to all positions on the block size limit are various images, conceptions, and assessments about the pros and cons of on-chain versus off-chain bitcoin transacting at both individual and systemic levels. My views on this topic have also evolved since my comments in 2016. I have become less skeptical about the potential for off-chain transacting, both conventional and cryptographic.

If bitcoin ever came to function as a common unit for the denomination of prices and payments, conventional payment systems could add it as an option with some back-end adjustments. This has not happened because 1) bitcoin is not a common unit for the denomination of prices and payments, providing little real demand for any such service, and 2) the prevailing climate of regulatory strangulation either dissuades or explicitly bans conventional companies from beginning to offer such a service in many cases.

This makes it essential to Bitcoin's practical viability that it has its own built-in independent payment infrastructure. Without one, it could not endure up to a point where it could become a unit for the denomination of prices and payments in broad enough demand that conventional payment services might support it and governments might grudgingly acquiesce to an overwhelming force of innovative success—as they have had to do in many past examples of world-changing technologies and practices that they had formerly long opposed and obstructed. Cryptographic off-chain solutions such as the Lightning Network are quasi-native bitcoin transfer technologies that stand in a middle ground on this spectrum between on-chain and conventional methods.

My formerly higher skepticism about the model of widespread off-chain transacting backed by on-chain bitcoin was based on two main factors. First, I have emphasized the model of unit/system duality in understanding the monetary valuation of bitcoin. In this view, the bitcoin unit is but a manifestation or aspect of the Bitcoin system. Part of the

unique valuation of bitcoin derives from the inherent inseparability of bitcoin/Bitcoin, unit and system. Second, in past commodity money systems, the model of backed substitute units has consistently degraded in time, eventually leading to systems of unbacked floating units known as fiat money.

On the first point, unit/system duality does not mean that, as a practical matter, substitute units cannot *also* be used as tokens of the underlying base unit. For example, a significant, easy-to-understand, and long-established method of off-chain bitcoin transacting has been exchange-trading. Trading happens on special-purpose trading engines while on-chain transactions are limited to customer deposits and withdrawals. The number, speed, and frequency of off-chain transactions far eclipse the related on-chain transactions, which are only needed to enter and exit the forum. What happens on the exchange could not happen on-chain, not even close. Participants in the exchange join it for convenience, speed, and liquidity, trading direct control of bitcoin for internal account credits. Clearing arrangements among exchanges could also be a route for bridging into more broadly usable payment systems should bitcoin ever become a common unit for the denomination of prices and payments.

It is important in monitoring off-chain developments to keep in mind that off-chain units are not bitcoin, but function as bitcoin *substitutes*. The only bitcoin units *are* on-chain units, much as the only gold coins *are* gold coins. Paper promissory notes and custodial account entries are *not*. Nevertheless, substitutes can also be used. In a commodity-money context, the most successful money substitutes have been described as *perfect* substitutes, which have "no difference between their value and that of the sum of money to which they referred…and…could not be subjected to an independent process of valuation on the part of those who dealt with them (Mises [1912] 1953, 74)." In the simplest terms, this would imply that a typical payee on the open market would not care whether they received a perfect substitute or the real thing (an off-chain unit or an on-chain one), though such a status is not necessary for all applications.

Many high-volume payment technologies already exist and are in wide use. They are just currently used to transfer conventional fiat currencies instead of bitcoin. The simplest long-term model to entertain would be conventional payment methods adding bitcoin to their currency options. This eventuality would not hinge on the widespread success of any particular cryptographic Layer 2 method. Yet such methods *might also* gain traction and compete against conventional options, further diversifying competing transfer methods. Direct on-chain transacting would also remain. The block size limit restricts the growth of on-chain traffic and chain size and promotes off-chain transacting methods, but a wide variety of competing off-chain technologies and services can facilitate the transfer of effective control of fungible bitcoin among parties.

On the second point, concerning problems and risks inherent to backing systems, Bitcoin the money unit has a unique advantage relative to previous commodity money backing systems in that its block chain is a readily auditable and universally accessible real-time public record that enables institutions to provide cryptographic proof of reserves. Vault auditing of gold reserves is naturally more opaque, expensive, and inaccessible, making metal reserve systems more susceptible to the kinds of encroaching corruption witnessed throughout monetary history. The idea of a "gold-backed cryptocurrency" thus omits one of bitcoin's headline features, that the monetary

commodity *itself*, not merely a substitute for it, "can be transported over a communications channel (Nakamoto 27 Aug 2010)."

As a bonus, running off-chain services normally entails operating a full node on the main network. A Layer 2 service can thereby fund the operation of Layer 1 nodes as an operating cost. Services can directly charge for traffic on Layer 2, which they have not been able to do on Layer 1. This also helps mitigate partial commons effects on Layer 1 full-node count. The "Bitcoin: Be your own bank" slogan has been used to mean that each individual connects to Layer 1 and operates there as his own personal bank. This could partly give way to "Bitcoin: Be a bank," meaning that any entity is free to connect a node to Layer 1, not only for itself, but also to have a go at being a Layer 2 service to customers. If one maintains that Bitcoin's main value is disintermediation, such a "re-intermediated" vision would be most unwelcome. However, if one views disintermediation as valuable but still secondary to the system's more critical and unique contribution of supporting a sound money unit, then this looks like one natural shape for a possibly successful long-term outcome (Graf 1 Oct 2019).

I classify bitcoin as a *digital monetary commodity* (with "monetary" as in *monetary asset* rather than "money") and a potential future *digital commodity money* (Graf 2015a). In contrast, *bitcoin substitutes,* unlike bitcoin, are in the strict economic-theory sense *tokens* in that they are meant to exchange in a fixed relationship with an underlying unit to temporarily represent it, much as a set of 100 pennies are valued as they are because one hundred of them are exchangeable for one dollar. They are defined as $1/100^{th}$ of a dollar. If money substitutes become weak of reputation, creating an expectation that they might not be freely exchangeable at promised rates, they can also be discounted, making them *imperfect substitutes*.

Note that other cryptocurrencies may be deemed competitors or attempts at offering "substitute" products on the market (for example, "use Litecoin instead of Bitcoin"), but this is a different sense than in the term *money substitute*, which has a far narrower meaning than *substitute product*. *Bitcoin substitutes* are therefore units that are **denominated in** on-chain bitcoin and circulate as temporary representations of such bitcoin for purposes such as convenience, speed, or privacy.

7. Applying evolutionary and ideological perspectives

The ongoing historical fact of the Bitcoin block size limit remaining in place cannot be understood only as the outcome of a rational decision or trade-off on the topic itself. Rather, it has come about bundled. One can say that the limit remains on the BTC chain in that the consensus has rejected a series of proposals to introduce a backward-incompatible hard fork to change it. Any such fork has its own separate and significant implications. A critical mass of relevant participants deemed avoiding one to be the best among imperfect practical trade-offs.

In economic-evolutionary metaphor, Bitcoin is a complex creature—closer to an ecosystem—managing to survive and expand in changing environmental niches. Evolutionary adaptation can only iterate on options already present or realistically buildable at each stage. A single factor cannot be targeted for change without far broader effects. Changing one factor is likely to mean changing a series of others, known or unknown, like it or not.

It is also important to note from an evolutionary standpoint that Bitcoin does not have to come to be used as a common unit for denominating prices and payments to play a valuable role. An ancient and still common error is to understand phenomena via a projected teleology, some imagined endpoint to which the observed thing is allegedly trending. The Darwinian revolution demonstrated a way to consider how existing systems *are* and how they *change* without dependence on any magnetic attraction from the future. A false view of evolution is encapsulated in the future-looking phrase "evolving toward." Instead of this, each step of evolution has its own current explanation and origins independent of any imagined future. Evolutionary understanding says: "this is how things are so far and how they have come about." While bitcoin **could** or **might** become a *digital commodity money*, its current and likely future role as a *digital monetary commodity* exists now and is comprehensible independently of any such future development.

A number of authors have used biological metaphors for Bitcoin. One effort to describe Bitcoin in terms of the fundamental characteristics of life included the following characterization of Bitcoin's block chain:

> Bitcoin…takes energy from the environment and puts things in order, i.e. it decreases its internal entropy. It does so by appending blocks to a well-ordered structure…This structure is just one part of a large and complex system, just like the backbone in vertebrates. It is important, no doubt. But distributed or not, a ledger on its own is as useful and as alive as a bag of bones. (Gigi 7 Aug 2019)

This leads us to a key argument in favor of hard-fork avoidance. Cultural and economic evolution proceed through memes, units of selection of ideas, rather than genes (for a masterful up-to-date treatment of *memetics* and its relationship to substantive content see Stewart-Williams 2018). A significant category of such memes is understood under the term *ideology*. Kurokawa (21 Aug 2018) explained the relevance of ideology to the context of Bitcoin and hard forks:

> The core foundation of any large group of people rests on ideology. Nations, religions, and political movements cannot exist without ideology and neither can cryptocurrencies. Stable ideologies allow communities to thrive…Bitcoin maximalists often say that the block size debate is not about the block size at all…The most important belief that the maximalists wanted to stand by in the block size debate is that backwards compatibility is never broken (or that we never hard fork). This may sound like a rigid requirement for a software project, but Bitcoin is not just a software project. It is a method of coordination for a large group of people who face extremely hostile and powerful adversaries. Understanding this fact, it becomes clear that software upgrades can be a large attack vector and may not be feasible when the adversaries are fully engaged.
>
> Critics are correct in saying that currently, the state level adversaries are not fully engaged and that hard forks are completely possible in practice. What they don't understand is the nature of ideology. Ideology can only be strengthened through strict adherence to it. A cryptocurrency project will not be able to easily switch to a policy of having no hard forks when the adversaries become suddenly engaged…Bitcoin users, who have been conditioned to believe that all hard forks are unsafe, will be immune when such an attack comes.

An ideology of hard-fork avoidance thus operates as a defense against potential assault. Provided the system is deemed to work 'well enough' for its most important purposes—which presumes a certain scope of common understanding of what those are (Section 6)—a high and multifaceted value attaches to proceeding with development in ways that avoid a hard fork. This goes beyond any single issue such as the block size limit being of this or that height, and also beyond the apparent and immediate trade-offs of hard forking versus other forms of network software revision viewed as competing ways to achieve coding objectives. Quasi-religious hard-fork avoidance is an ideological rule-of-thumb that promotes the long-term resilience of the system and its community.

In bundled Darwinian fashion, maintaining a hard-fork avoidance ideology would also operate to protect the monetary policy. Even if many actors behind the consensus were to become attracted to an inflationary direction, such a shift could not be executed without *also* going against the hard-fork avoidance principle. This places one additional layer of protection over the ongoing monetary wisdom of the human consensus.

Not only has hard-fork avoidance come to be supported by Bitcoin (BTC) community ideology, it must also be understood as one structural tendency of Bitcoin's governance model. With a backward-compatible soft fork, new software features can be introduced and made available while anyone running other consensus-code compatible software versions can continue as before and ignore the new features. This is not the case with a hard fork. If a new hard-fork version gains wide adoption, any user still running an older version will be rendered incompatible. This expands the broadness of agreement needed before proceeding with low risk of a chain split.

This perspective on the ongoing role of ideology may appear to contrast with the view that "Bitcoin is sovereign," that its core elements cannot be altered. An emphasis on ideology would seem to imply that mere shifts in sentiment could in fact lead to core rules being altered. The reality is that if the shift were broad enough, they could be. However, maintenance of an ideological sentiment that Bitcoin's consensus rules are unalterable helps to make this the case in a self-fulfilling fashion. A contrasting ideology that hard forks are easy and consensus rules can be changed is likewise more likely to lead to them being changed. Although Bitcoin's structural dynamics have strong influences through incentives, consciousness and culture also remain active in parallel. We are creatures not only of choice, but also habit and precedent.

Contrasting case studies in hard-fork-embracing ideology

Some empirical evidence from the past few years appears to support the evolutionary adaptiveness of hard-fork avoidance as an ideological strategy for decentralized digital cash. The Bitcoin Cash (BCH) chain split of 1 August 2017 created a hard fork of Bitcoin (BTC) with an increased block size limit. Since both chains continued from the last common-ancestor block #478558, the result is described as a chain split.

The community of Bitcoin Cash developers not only chose a hard fork to make a single change, but also consciously departed from hard-fork avoidance. They set a course of embracing hard forks and using them on a routine basis as a normal method of making software revisions in a way perceived to be more efficient.

As I wrote immediately after the Bitcoin Cash chain split, if the split was to be an experimental "test" of differences in block size limit, it would be a poor and confounded one because so many other variables also contrasted, including: "the presence/absence of SegWit, the respective quality levels and reputations of software development teams and software testing processes, differences in user traffic, and the extent and stability of relative hashing power (Graf 5 Aug 2017)."

To this list of confounding variables, I add differences in ideological orientation toward hard forks. Any consensus shift to hard fork BTC requires a very high, and possibly rising, standard of evidence, and cannot be viewed only as one more practical option among software revision methods, as it has become in competing projects.

Many critics at the time also considered it a significant negative that the BCH movement was led by a relatively small collection of prominent individuals from within wider Bitcoin mining, development, and business communities. This would tend to make the BCH project more malleable/flexible compared to BTC. Whether such changeability was essentially positive or negative was itself a subject of disagreement.

After two years in operation, Bitcoin Cash has attracted little of Bitcoin's network effect, price, or hashing power. BCH typically trades at just under 3% the price of BTC. Moreover, the hard-fork-embracing ideology seems to have returned to bite the project. On 15 November 2018, 15.5 months after the BCH chain split, Bitcoin Satoshi's Vision (BSV), was, in turn, launched in a hard-fork chain split from BCH. The hard-forking chickens had come home to roost.

But whereas the immediate aftermath of the BTC/BCH split saw the combined value of both coins rise, the BCH/BSV split was less "orderly," and the BCH price declined and never fully recovered. BCH had traded at 6–7% of BTC prior to the BSV split but has traded at 2–4% after it. BSV, for its part has traded at about 1–2% of BTC since its launch. In other words, the value of BCH+BSV has not recovered to the pre-split value of BCH alone in BTC terms. Repetitive splitting of chains with attendant multiplication of unit types is anathema to a hard-money application.

Ethereum, though intended for purposes other than decentralized digital cash, likewise went through a chain-splitting hard fork. This left two separate chains starting on 20 July 2016, Ethereum (ETH), which introduced changes to reverse data that had been recorded on its block chain, and Ethereum Classic (ETC), which rejected those reversals and carried on the original chain. The chain-data conserving ETC trades at about 3% of the chain-data altering ETH, the inverse of the BTC/BCH price relationship.

It is notable in this connection that Ethereum has an identifiable active founding group and foundation that benefitted from a 72mn coin pre-mine and which still exercise considerable influence on development and development funding. The decision to hard fork to reverse chain data was conducted through voting based on coin ownership (Madeira 2019). An identifiable, active founding group—of the kind that Bitcoin uniquely lacks—could play a role in altering chain data and a collective voting method was used in contrast to Bitcoin's emergent consensus.

Both the BTC/BCH split and the ETH/ETC split have produced starkly asymmetric outcomes in terms of price and hash rate within the two pairings. As speculation on a

complex series of unique events, ETH's far higher degree of leadership- and financial-centralization as compared with BTC may have helped enable it to tip the scales of the network effect toward its own chain-altering hard fork split, whereas BTC maintained network-effect dominance as the non-hard-forking side in the BCH chain-split event.

Even though Ethereum Classic was created so as not to acquiesce in a hard fork that would change data written to an "immutable" block chain, the ETC project nevertheless engages in hard forks for software revision. Bitcoin thus uniquely stands out from among these four projects as representing hard-fork avoidance, rejecting hard-forking even as a means of software revision, let alone chain-data revision.

Argumentation, rhetoric, and ideological drift can enhance or reduce the likelihood of a hard fork. Yet hard-fork avoidance must also be understood as having a structural foundation in Bitcoin's consensus context, uniquely devoid of an active founding figure or group with a special position from which to lead and encourage a hard fork. Incentive-based structural tendencies and specific cultural and ideological components operate in parallel and are not mutually exclusive. Bitcoin-community ideology has come to support a climate in which, "if anything is even vaguely debatable, then nothing will happen (Vays, 17 Aug 2019)," which reinforces the system's incentive-based tendencies. For certain peculiar and unique applications—such as a global digital hard-money unit—high risk-aversion in the supporting public club might be part of what is required to survive and perhaps eventually thrive in such a role.

8. Concluding summary

Examination of the Bitcoin block size limit issue requires placing it within several layers of wider context. The economic analysis of the policy itself can identify sets of pros and cons likely to follow from it. However, the policy is embedded in a context that renders it highly interdependent with other issues, foremost among which are the wider implications of hard-fork avoidance versus hard-fork embrace.

Also critical is the analysis of the relationship between money and money substitutes and views on the viability of various methods for issuing bitcoin substitutes, especially should bitcoin, now a *digital monetary commodity*, come to be a unit used in the denomination of prices and payments, making it a *digital commodity money.* Even though bitcoin's value is best understood in terms of its dualistic inseparability from the Bitcoin system, the unit/system-duality value model does not imply that substitute tokens cannot *also* be usefully employed, valued as such.

The functions of payment making and the determinants of the economic value of monetary units are distinct topics, with separate sets of competitive factors. Existing payment systems, for example, easily operate in various currencies depending on the locations and preferences of their users around the world. At a purely technical level, adding cryptocurrency-denominated payment options via conventional technologies might be accomplished through back-end modifications to existing systems, should demand expand.

The block size limit represents a form of *voluntary-sector artificial scarcity*, and as such is not subject to legal-theory critiques of compulsory-sector artificial scarcity measures such as IP laws. The limit operates as one among possible means of addressing apparent

tragedy-of-the-commons dynamics inherent to the novel combination of a shell of natural scarcity protecting a core of artificial scarcity, specifically, the tendency for traffic flow and accumulated chain size in the artificially scarce core to build without direct pricing limits, increasing future participation costs in the naturally scarce shell (largely computational capacity and electricity).

While the naturally scarce shell is subject to the usual analyses of property rights and supply and demand, the artificially scarce core—what is being protected—requires symbolic-realm measures for its competitive structuring. These artificial-scarcity measures along with their cost/benefit packages include *club goods* such as the block chain that are non-rivalrous in being shared by users as a common resource and partly excludable in that consensus rules place general limiting conditions on chain additions.

Property rights can alleviate commons problems in ownable resources by internalizing their future discounted value to specific decisionmakers. However, property rights cannot be validly applied to the purely symbolic realm. Analogous commons problems in unownable resources may be addressed instead through private rules and private governance. If the rate of block chain size growth is viewed as such a commons problem, club governance rules impacting participants are the natural venue for its regulation.

The block size limit restricts the growth of on-chain transacting and promotes off-chain methods of all kinds, backed with on-chain units and settled in on-chain transfers. Control of bitcoin units can be transferred among parties not only on chain, but also through a wide variety of competing technologies and services. Bitcoin has a fresh advantage relative to previous commodity money backing systems in that it is a readily auditable and universally accessible real-time public record that enables institutions to provide cryptographic proof of reserves.

Bitcoin participation and exclusion are automatic based on compliance with consensus code. Anyone who follows the rules is free to participate without any human-mediated procedure, making it unlike private clubs. Nevertheless, just as with Stringham's *private governance* model, the actual net pros and cons of operative rule sets lead to an ongoing Hayekian discovery process of comparative successes. Key to this process, the net costs and benefits of rule sets tend to be internalized to the particular project purveying them and its participants, allowing a "competition in rules" to identify rule sets with apparently superior net benefits in practice and not only in theory.

Bitcoin's *code-enhanced governance* renders any changes to its consensus code extremely difficult to implement, reinforcing a status quo bias within the open *public club* of opt-in participants. This not only aids in protecting the monetary policy, but also provides the block size limit with a memetic evolutionary survival advantage simply in that it is already in place. In the context of decentralized digital cash, 1) the absence of an active dominant founding figure, group, or foundation 2) hard-fork avoidance ideology, and 3) a fixed block size limit, have thus far been associated with overwhelming economic dominance in open voluntary-sector competition. This is consistent with the view that Bitcoin's primary competitive advantage and most unique and valuable contribution lie in the credibility of its ongoing maintenance and enforcement of a strictly limited unit production schedule.

References

Ammous, Saifedean. 2018. *The Bitcoin Standard: The Decentralized Alternative to Central Banking.* Hoboken, New Jersey: John Wiley & Co.

Buchanan, James M. 1965. "An Economic Theory of Clubs." *Economica*. New Series, Vol. 32, No. 125, 1–14.

Graf, Konrad S. 24 Oct 2014. "Sidechained bitcoin substitutes: A monetary commentary." *konradsgraf.com.*

———. 2015. *Are Bitcoins Ownable? Property Rights, IP Wrongs, and Legal Theory Implications*. CreateSpace Publishing.

———. 2015a. "Commodity, scarcity, and monetary value theory in light of bitcoin." *The Journal of Prices and Markets*, Volume 3, Issue 3 (Winter).

———. 4 May 2016. "Bitcoin block size political economy." *Bitcoin.com.* PDF on *konradsgraf.com*.

———. 8 Jul 2016. "Software choice, market differentiation, and term selection." *konradsgraf.com*.

———. 9 Jul 2016. "Market intervention through voluntary community rules." *konradsgraf.com*.

———. 10 Jul 2016. "Differentiation from the 21-million coin production schedule." *konradsgraf.com*.

———. 5 Aug 2017. "Descendants with modifications: Bitcoin's new and possibly beneficial evolutionary test." *konradsgraf.com.*

———. 1 Oct 2019. "Sound money strikes at the root: A review essay on *The Bitcoin Standard*. konradsgraf.com.

Gigi. 7 Aug 2019. "Proof of life: Why Bitcoin is a living organism." *Medium.com.*

Hoppe, Hans-Hermann. 2010 [1989]. *A Theory of Socialism and Capitalism.* Auburn, Alabama: Mises Institute.

Kurokawa, Kay. 21 Aug 2018. "The Bitcoin Cash ideology and the incoming schism." *Hackernoon.com.*

Madeira, Antonio. 12 Mar 2019. "The Dao, the hack, the soft fork and the hard fork." *Cryptocompare.com.*

Mises, Ludwig von. 1953 [1912]. *The Theory of Money and* Credit. New Haven: Yale University Press. mises.org/document/194/The-Theory-of-Money-and-Credit

Nakamoto, Satoshi. 15 Nov 2008. "Bitcoin P2P e-cash paper" entry #014858. satoshi.nakamotoinstitute.org/emails/cryptography/threads/1/#014858

———. 27 Aug 2010. "Re: Bitcoin does NOT violate Mises' Regression Theorem," entry #11405. bitcointalk.org/index.php?topic=583.msg11405#msg11405

Stringham, Edward. 2015. *Private Governance: Creating Order in Economic and Social Life*. Oxford University Press.

Stewart-Williams, Steve. 2018. *The Ape that Understood the Universe: How the Mind and Culture Evolve.* Cambridge University Press.

Vays, Tone. 17 Aug 2019. "Bitcoin: There can be only one" [presentation transcript]. diyhpl.us/wiki/transcripts/bit-block-boom/2019/there-can-only-be-one/