

An (Institutional) Investor's Take on Cryptoassets

December 24, 2017 • version 6¹

John Pfeffer

[Medium](#) • [@jlppfeffer](#) • [LinkedIn](#)

John Pfeffer is an entrepreneur and investor. In the 2000s, he was a London-based partner at private equity firm Kohlberg Kravis Roberts, and in the 1990s, he was Chairman of the Executive Board of leading French IT company Groupe Allium S.A. Before that, he advised on turnarounds while with McKinsey in Europe and Latin America.

IMPORTANT NOTICE: *This document is intended for informational purposes only. The views expressed in this document are not, and should not be construed as, investment advice or recommendations. Recipients of this document should do their own due diligence, taking into account their specific financial circumstances, investment objectives and risk tolerance (which are not considered in this document) before investing. This document is not an offer, nor the solicitation of an offer, to buy or sell any of the assets mentioned herein.*

Amidst the indiscriminate speculation, sensationalist and mostly misguided media coverage and roller-coaster price volatility, this paper sets out to consider cryptoassets from the perspective of a rational, long-term investor. As investors, we look for things that generate sustainable, ideally growing economic rent—an economic surplus that will accrete to us. This paper evaluates the extent to which cryptoassets offer the foregoing. It aims to assess the potential future value of cryptoassets at mature equilibrium,² on the assumption that they *develop successfully and achieve widescale adoption*. By design, it does not dwell on the significant risks that a given cryptoasset could fail, for technical, regulatory, political, or other reasons. These risks are very real, and are well documented elsewhere. Temporarily setting them aside allows for an objective analysis of the *potential* value of different kinds of cryptoassets and their use cases.

I write not from the perspective of a trader, but from that of an investor who believes the long term is easier to predict than the short term. The paper thus focuses entirely on long-term equilibrium outcomes and investment strategy rather than short-term price movements. It also assumes the reader has some familiarity with the topic.

Blockchain technology has the potential to disrupt a number of industries and to create significant economic surplus. The open-source nature of public blockchain protocols,

¹ Earlier versions of this paper were drafted beginning in June 2017.

² The notion of mature equilibrium as I use it here is admittedly imprecise. Conceptually I mean once the speculative phase has passed and (i) in the case of monetary store of value, once there is a mainstream, institutional view that crypto is a core monetary store of value like gold is today and (ii) in the context of infrastructure and applications, once markets are valuing cryptoassets based on significant realised user penetration. The obvious analogy is the internet. Internet penetration and internet-enabled businesses are still growing today but growth is slowing. Today, large internet-enabled businesses are valued based on financial ratios such as PEG and EBITDA multiples rather than clicks or eyeballs as was the case in the late 1990s. That's the end point I'm thinking about. For shorthand, let's assume 10 years from now.

combined with intrinsic mechanisms to break down monopoly effects, mean that the vast majority of this economic surplus will accrue to users. While tens or perhaps hundreds of billions of dollars of value will also likely accrue to the cryptoassets underlying these protocols and therefore to investors in them, this potential value will be fragmented across many different protocols and is generally insufficient in relation to current valuations to offer a long-term investor attractive returns relative to the inherent risks. The one key exception is the potential for a cryptoasset to emerge as a dominant, non-sovereign monetary store of value, which could be worth many trillions of dollars. While also risky, this potential value and the probability that it might develop for the current leading candidate for this use case (Bitcoin) would appear to be sufficiently high to make it rational for many investors to allocate a small portion of their assets to Bitcoin with a long-term investment horizon.

We can break cryptographic token use cases into three broad categories:

1. Network backbone / Virtual Machine (e.g., Ethereum)
2. Distributed applications (Dapps)
3. Money, and in particular:
 - a. Payments
 - b. Monetary store of value.

I will start by looking at the first two use cases from a general perspective and then dive deeper in analysing the largest current example of the first one, Ethereum. I'll then turn to a discussion of the different functions of money, the potential for cryptoassets to perform them and the implications for the value of such cryptoassets, including Bitcoin.

The economics and valuation of utility protocols

Use cases 1 and 2 can be grouped into what I call utility protocols. I will start with some general observations on utility protocols and the implications for their network valuation at equilibrium and then specifically consider the network value of Ethereum at mature equilibrium.

General observations

A blockchain protocol is a database maintained by a decentralised consensus mechanism operated by its nodes. Utility protocol tokens serve to provision scarce network resources: the processing power, memory, and bandwidth necessary for maintaining the blockchain in question. These resources have a real-world cost in terms of energy and the equipment employed, and these costs are borne by the miners who maintain the blockchain by providing computational services. The miners may be remunerated for their service with block rewards, paid in protocol tokens, and/or transaction fees, paid in protocol tokens or some other means of exchange. While protocol developers may claim that tokens are the basis for other kinds of exchange among users and not just a means of allocating and paying for computing resources, it is my argument that, at mature equilibrium, tokens will do no more than allocate computing resource, with the exception of the special case of a cryptoasset that serves as a monetary store of value.

A given protocol is analogous to a simplified economy. The GDP of such an economy would be the aggregate cost of the computing resources necessary to maintain the blockchain, based on the quantity of processing power, memory and bandwidth consumed, multiplied by the unit cost of each. The token is typically the currency used to pay for those resources. The total network value is analogous to the money supply M (i.e., all tokens in issuance), where $M = PQ/V$; PQ (Price x Quantity) is the total cost of the computing resources consumed, V is

a measure of how frequently a token is used and reused in the system (its velocity, V). The value of a single token is therefore M/T , where T is the total number of tokens.

If a given utility protocol does not have a built-in mechanism, such as Ethereum's GASPRICE, to ensure that the cost of using the network does not arbitrarily and sustainably diverge from the underlying cost of the computational resources it consumes, one of three things happens: (a) the token's price trades to a level such that there is no premium cost to using the network (i.e., there is no economic rent); (b) the chain forks into a functionally identical but less rent-seeking chains until any premium usage cost and economic rent on the network declines to a level at which it is no longer worthwhile to arbitrage; (c) the protocol's adoption is temporarily limited to the highest-value use cases until (a) or (b) occurs. In all cases, the equilibrium result must be at or near marginal revenue = marginal cost for the mining industry maintaining the blockchain in question, so that the token's value cannot materially decouple from the underlying computing resource cost.

PQ, the cost of computing resources required to maintain a blockchain, is not only low relative to the current network values being attributed to cryptoassets; it is also inflated by the prevalence of proof-of-work consensus mechanisms, which mean that the vast majority of computing resource consumed is make-work. To the extent that new scaling technologies such as proof-of-stake, sharding, Segregated Witness, Lightning, Raiden and Plasma become prevalent, the amount of computing resource consumed may become quite small. Note also that in the context of cryptoassets, V could go very high at equilibrium. Even if a significant portion of a given cryptoasset has a low velocity because it is being hodl'd by speculators or because it is staked by miners under a proof-of-stake consensus mechanism, the circulating portion of the tokens can circulate at the speed of computer processing and bandwidth—i.e., fast and accelerating. The implication is that average velocities can and are likely to be high, regardless of how many tokens are actually actively circulating for utility purposes to allocate network resources.³ The combined effect of low and falling PQ and potentially very high V is that the utility value of utility cryptoassets at equilibrium should in fact be relatively low.

Clearly, scaling solutions such as proof-of-stake, etc. are bullish for adoption/users but bearish for token value/investors. Even without those technology shifts, the cost of using decentralised protocols is deflationary, since the cost of processing power, storage and bandwidth are deflationary. This is also bullish for adoption and users and bearish for token value and investors.⁴

Whatever scaling solutions are developed, the inherent redundancy of the consensus mechanism means that there may be fewer use cases than many decentralised revolutionaries think in which a decentralised solution displaces a centralised solution. Use cases will be limited to dematerialised networks where the value of decentralisation, censorship-resistance and trustlessness is high enough to justify the inherent inefficiency and redundancy of the consensus mechanism. Is it worth the cost for payments? Yes for some, but not for all. Consider Twitter -- what is the added value to the user of a massively redundant, trustless,

³ Chris Burniske's recent blog post "Cryptoasset Valuations" (<https://medium.com/@cburniske/cryptoasset-valuations-ac83479ffca7>) estimates an average V of 7, after adjusting for hodlers, stakers, etc. This assumption may be optimistic (meaning, it is probably a low value of V to assume at equilibrium and therefore an optimistic number to be using to estimate the potential equilibrium value of a given cryptoasset), but his framework is useful for thinking about the different drivers of V for a given cryptoasset.

⁴ I have yet to come across any examples of a protocol where I have been persuaded that when all is said and done the underlying scarce resource being provisioned is something other than computing resources, or at least where that is what it will boil down to at competitive equilibrium after competition in mining, forks, etc. Please alert me to any counter examples you have seen or can think of.

decentralised Twitter? Is that added value enough to offset its inefficiency compared to the incumbent centralised Twitter? Would Token Twitter offer compellingly higher utility compared to centralised Twitter, including enough surplus utility to offset the cost of operating the consensus mechanism? I'm not so sure.

People often make the mistake of conflating the monopoly network effects of, say, Facebook to blockchain protocols. This notion is fallacious on several levels:

- Blockchain protocols can be forked to a functionally identical blockchain with the same history and users up to that moment if a parent chain persists in being arbitrarily expensive to use (i.e., rent-seeking). Like TCP/IP but unlike Facebook, blockchain protocols are open-source software that anyone can copy or fork freely. A protocol fork is analogous to a team of Facebook developers who decide one Tuesday morning that Zuck is not paying them enough; they could simply flip a switch and use the servers and software that run Facebook to run a new Facebook that is functionally identical, with all the same users and data up to that point. That can, does, and will happen all the time in protocol-land, but would be theft in the context of private companies that own their code, data, intellectual property, etc. Those property rights are why Zuck is rich, and their absence in the protocol economy has profound implications. The ability to fork protocols maximises utility for users but suppresses economic rent for token holders.
- When people talk about the potential value of cryptoassets, they often refer to Metcalfe's Law. Metcalfe's Law asserts that network value = $\Theta \cdot n(n-1)$, where Θ is a constant that captures the differences in the economics built into the business model of each network and where n is the number of nodes in the network. It's not enough to focus on $n(n-1)$. You must also consider what Θ is. Wikipedia has a lot of contributors and users but not a lot of monetary value because it doesn't charge users or have advertisers or attract any other sources of revenue apart from donations. Facebook's Θ is higher than Twitter's because its advertising business model is stronger. TCP/IP lacks financial value not only because no one owns it but also because it doesn't have a revenue model. The problem for utility protocols is that the Θ in question is driven by the cost of the computing resources to maintain the network, which is relatively low and deflationary and which must remain low for their adoption to be successful vs. non-distributed technologies.
- When thinking about whether a protocol's token can capture and sustain economic rent, what is relevant is whether the mining industry maintaining the protocol's blockchain is competitive, not the stickiness of users. The mining industry supporting any decentralised protocol must be a competitive market; otherwise the protocol isn't decentralised. It is the economic competition amongst miners that will ultimately drive the cost of using the protocol and therefore the value of the token. No mechanisms for monopoly rents there.
- Not only must protocols compete against their own potential forks; competition amongst protocols is also fierce. Witness, for example, recent press reports that Kik is considering migrating its token network from the Ethereum backbone to another blockchain because the Ethereum network is becoming too expensive to use.⁵

⁵ <https://www.coindesk.com/kik-might-move-its-ico-tokens-to-a-new-blockchain/>

- The network value of a tokenised version of a dematerialised network business (a social network, Uber, AirBnB, a betting exchange, etc.) will by construction be a small fraction of the enterprise value of its centralised, joint-stock-company equivalent. Holding the number of users constant, you basically take the fully-loaded IT budget (including energy and a capital charge) of those companies (representing PQ) and divide by some (likely high) velocity V. The disruption of traditional networked businesses by decentralised protocol challengers will represent an enormous transfer of utility to users and an enormous destruction of market value. Great for users, the economy and society; bad for investors.

The next topic to address is the impact of a move to proof-of-stake mining and of staking models in general on the network values of Ethereum and other protocols. The idea is that miners are compensated for maintaining the network either in a native cryptoasset or another cryptoasset (such as ETH or BTC), in proportion to the amount of the native network cryptoasset that they stake (i.e., effectively put into escrow and at risk of loss if they attempt to validate false transactions and the like). The promoters of this idea hope that it will reduce the actual computing costs of maintaining the network, by eliminating the costly proof-of-work mechanism, while at the same time creating an alchemic virtuous cycle wherein miners buy and lock up significant amounts of the native cryptoasset as an investment conveying them a right to a mining revenue stream, thereby reducing the velocity of the native cryptoasset and causing its value to rise to a level representing some multiple of their mining profits, much as taxi medallions or shares in a company are valued based on the net present value of future cash flows.

Let's think through how this plays out.

First, before staking is introduced into the equation, we've established that forks and competition in mining and among protocols lead us to an equilibrium outcome where PQ equals the aggregate cost of the computational resources (capital charge on or usage cost of processing and storage hardware, cost of bandwidth and energy) of maintaining the network.

Second, recall that the impetus for moving from proof-of-work to proof-of-stake is to reduce the amount of computational resource and energy required to maintain the network by a couple orders of magnitude. That's good for scalability and potential adoption, but also means a commensurate reduction in the PQ of the network.

Third, let's layer on the idea that in order to participate in mining and the associated revenues, on top of paying for processing power, storage, bandwidth and energy, you must now bear an additional cost in the form of a capital charge from acquiring and immobilising an amount of the native cryptoasset. This capital charge on immobilised cryptoasset is added to PQ, making the protocol in question more expensive to use than an equivalent utility protocol that doesn't require staking (or where staking is less expensive because the native cryptoasset is cheaper).

This system operates a bit like a taxi medallion system: an authority issues a finite number of licenses, and you must buy one from another medallion holder if you want to operate a taxi. The value of the license captures the discounted value of any economic profits that are expected to accrue from operation. Whoever owned the license first is the primary beneficiary of this monopoly, and he receives that value when he sells the license to someone. The buyer of the license does not enjoy any economic rent because he paid the discounted present value of it to the previous license holder, and so on as the license changes hands. Passengers pay higher fares because the taxi driver's capital cost of buying the license must be compensated for, all for the benefit of the first owner of the license.

Imagine there are several different taxi companies operating that have acquired a number of licenses. Now imagine that a new entrant decides it would like to take market share. In the world of a taxi medallion monopoly created by an issuing public authority, they would have no option other than to buy medallions from other medallion owners. But here is where protocol-land is different from real-world taxi medallion schemes. Protocols are open source software and can be freely forked.

In protocol-land, all the upstart taxi company needs to do is to fork the protocol, effectively issuing an identical number of new taxi medallions and reallocating medallions owned by existing large taxi companies to itself and perhaps a few other friends. Because the upstart taxi company didn't have to pay for its taxi medallions, it and the other recipients of the new medallions can charge its passengers lower fares. Passengers thus flock to the upstart company, and the monopoly value embedded in the original taxi medallions vanishes. Everyone in the system except for the large taxi company wins. If necessary, this process can be repeated indefinitely. The result is that the medallions have low values (as would the analogous native cryptoasset).⁶

Another mechanism for utility protocols is mine and burn. In this system, new coins are minted and allocated to miners based on the network services they provide, and users must buy these coins and burn them to pay for transaction processing. This is a perfectly fine mechanism, but it simply ensures that the network value equals PQ/V , where PQ is the actual fiat cost of maintaining the network and V is the average time from minting to burning. That gets you to the same low equilibrium network value more simply and quickly.

Other general observations:

- Analysts often use a working capital analogy in order to assess how much of a given cryptoasset a user will stock to facilitate actual use of a given blockchain's utility function. Fair enough, but digging further into that line of thinking, the way optimal inventories of a good are set is based on the relationship of the volume and volatility of demand, optimal order sizes, communication and delivery latency and production times. Since cryptoassets are generally highly divisible and may circulate very fast (as fast as processor speed and bandwidth allow), it would seem to me that a user would, by the same maths as those used to determine optimal inventory quantities, conclude that he needs to hold very little inventory of a given cryptoasset. Friction moving among cryptoassets is already low and will quickly disappear entirely with technologies like atomic swaps. Consequently, one would expect velocity to be very high at equilibrium. It would make no more sense for users to hoard utility cryptoassets beyond the minimum they need to carry out their desired operations than it would be for individuals to hoard petrol or for companies to hoard giant warehouses full of whatever goods they sell. Companies need inventories of goods to run a business and those inventories have a value on their balance sheet, but they try to minimise such holdings, as they are unproductive assets that are costly to finance and carry. They certainly don't try to accumulate more inventory than necessary as a way to store their retained earnings. Similarly, individuals have petrol in the tanks of their cars, but they don't stockpile petrol in their basements as a form of savings.⁷

⁶ The competitive forces to eliminate economic rent would function in largely the same way whether the staking system involves payment for services in cryptoassets that are native or external to the protocol at hand.

⁷ See also Vitalik Buterin's recent blog post: "On Medium of Exchange Token Valuations" (<http://vitalik.ca/general/2017/10/17/moe.html>)

- For every successful utility protocol (certainly for every successful Dapp), there will be n failed versions. In fact, one of the advantages of the protocol economy is that it facilitates open and inexpensive experimentation, which will mean that there will be many more attempts and many more failures, and that each success will be individually smaller in its value and reach. The open-source, forkable nature of this kind of software will likely drive toward a fragmentation of use cases and protocol functionality; businesses built on top of the protocols will be protocol agnostic and capable of using and combining modularly a changing array of protocols to deliver whatever service or value chain they are trying to deliver. These dynamics are great for users and generate lots of positive economic and social externalities, but they are bad dynamics for investors.⁸ The problem of making money by investing in utility protocols is aggravated by: (a) the fact that this is a fragmented space with very high failure rates, so selecting winners *a priori* will be very difficult; and (b) the fact that most of the long-term winning protocols probably haven't even been launched yet (witness the fact that the most valuable internet businesses were founded after 2001).
- Developer incentives over time are a fundamental issue in crypto. For most protocols, such incentives are heavily front-ended around launch and insufficiently provided for over time. The more ambitious and long-term a protocol's development roadmap is, the more problematic this failure of incentives becomes. The incentives to improve an existing protocol by forking it may be strong if some tokens are reallocated at the fork to the developers making the improvements. For example, where tokens have been retained by a foundation linked to the original protocol developers, an aggressive group of forking developers could reallocate the foundation's tokens to their own new entity in their fork, leaving all other users in the same position and letting the market decide which fork to support. The incentives for a developer to create a new, competing protocol are also strong, but network effects do make it harder to displace an existing protocol than to improve or fork it. Miners and perhaps large users have a strong economic incentive to invest in development of the protocol they are mining either through changes to the protocol or by forking it. The foregoing suggests that we're likely to see (a) more success with protocols focused on simple use-cases that require less ambitious future development; (b) future protocols launched with better long-term developer incentive schemes (easier said than done)⁹; (c) aggressive forks that transfer value from incumbent to challenger developers¹⁰; and (d) large miners/users or groups of miners/users acting together employing or paying developers to improve legacy protocols either directly or via forks.

The implication of this section is not that utility protocols won't have any network value. PQ/V does represent positive value. The implication is that network value of a utility

⁸ See also Teemu Paivenen's blog post "Thin Protocols" (<https://blog.zepelin.solutions/thin-protocols-cc872258379f>).

⁹ Tezos proposes an interesting potential solution to the developer incentive problem. Tezos combines a PoS consensus mechanism with a system whereby token-holders can vote on improvements to the protocol proposed by developers and reward the developers for their contribution. We'll see if it works, but the problem for Tezos remains that the mature equilibrium value of the Tezos token will be $T_{\text{Tezos}} = PQ/VM$ where PQ is the cost of the computing resource maintaining the Tezos blockchain, i.e., T_{Tezos} probably won't have a high value when the dust settles.

¹⁰ See also Fred Ersham's blogpost "Accelerating Evolution through Forking" (<https://medium.com/@FEhrsam/accelerating-evolution-through-forking-6b0bba85a2ba>)

protocol will converge on or near an equilibrium, where it is a fraction (denominator V) of the actual cost of the computing resources consumed to maintain the networks.

For a fork to succeed, there needs to be enough value available to arbitrage to incentivise users, some miners and a sufficiently credible developer group to support the fork. It should therefore be acknowledged that, to the extent the equilibrium outcome is arrived by way of one or more forks, there could be a sustainable level of network value economic rent premium above computing cost that is too small to provide adequate incentives for a fork to succeed. I would not, however, consider it to be a very compelling investment thesis when the best I can hope for is to keep an amount of value corresponding to an economic rent that's too small for anyone to bother arbitraging it away from me despite relatively low barriers to doing so. While a protocol's core development team may be bound by various soft ties, in protocol-land (unlike in a traditional software business), the work product is all open-source; intellectual property isn't generally owned or protected; and developers have little or nothing in the way of contractual ties or limitations (e.g., no non-compete, no non-disclosure, no non-solicit). That means developers can defect or take the work of others. At a minimum, these factors place a low ceiling on how much economic rent can be created and sustained.¹¹

As illustrated in the ETH valuation example to follow, it is likely that the combined network values of all utility protocol cryptoassets together will total between tens of billions and hundreds of billions of dollars. That is significant value, but not when compared to the current ~\$250 billion combined network value of protocols other than Bitcoin. Investing in utility protocol cryptoassets could make sense if their current network values were one or two orders of magnitude lower than they currently are, but at current valuations, the risk/return to investors is not attractive.

The Network Value of ETH

ETH, the Ethereum token, is an interesting case to explore because of its significant current network value and Ethereum's potential as the ultimate utility protocol. Ethereum could serve as the backbone for processing smart contract operations for (hopefully) untold numbers of decentralised applications, DAOs, etc., and perhaps one day maybe even something like the fabled Ethereum Virtual Machine (EVM).

Ethereum's developers understood that for Ethereum to fulfil its potential, the cost of using it as a smart-contract-executing utility must be as low as possible and must not depart at equilibrium from the actual cost of the computational resources consumed. To ensure this will be the case, they built the GAS mechanism into Ethereum to decouple the use of the network (and the cost thereof) from the value of the ETH token.

Each possible type of computing operation has a pre-defined GASCOST, measured in units of GAS. GAS may then be paid for using ETH (or another token or currency) based on the GASPRICE 'exchange rate', which is freely set among users and miners.¹²

¹¹ An interesting business idea that someone could logically pursue at some point would be to raise capital to fund a crack team of mercenary blockchain developers and systematically target technically-mature or maturing protocols where there is still a significant economic rent premium and arbitrage that value via hostile forks of those protocols in a way that reduces cost and/or improves functionality to users and reassigns network tokens held by the incumbent developer team and backers to the insurgent team and backers.

¹² Note that because GASPRICE is fully-flexible, GASCOST might only need to be updated in the system from time to time if and to the extent the relative cost of certain sub-components of computing costs changes, for example the cost of processing power vs storage.

The Ethereum Homestead Documentation makes this all clear:

“Gas Price is how much Gas costs in terms of another currency or token like Ether. To stabilise the value of gas, the Gas Price is a floating value such that if the cost of tokens or currency fluctuates, the Gas Price changes to keep the same real value. The Gas Price is set by the equilibrium price of how much users are willing to spend, and how much processing nodes are willing to accept¹³.” (Ethereum Homestead Documentation Release 0.1, p49)

“Gas and ether are decoupled deliberately since units of gas align with computation units having a natural cost, while the price of ether generally fluctuates as a result of market forces. The two are mediated by a free market: the price of gas is actually decided by the miners, who can refuse to process a transaction with a lower gas price than their minimum limit.” (Ethereum Homestead Documentation Release 0.1, p68)

This is all logical in the sense that GAS, and by extension the ETH token itself, is a metering device meant to ensure correct economic allocation and remuneration of the network’s resources. In the long term, the GASPRICE (and through it the value of ETH) should therefore tend toward the actual marginal cost of computing resource on the network. It could not possibly be otherwise, since if the cost of running operations on the Ethereum blockchain became materially more expensive than the actual underlying cost of computing resources consumed by it, people would simply use another blockchain where that premium doesn’t exist (or fork to create a cheaper Ethereum network that has identical functionality and users at that moment)? Also, if the GASPRICE were to decouple sustainably from the actual computing cost of operations, then mining would be the only perfectly competitive industry in history to earn sustainably positive economic rent. There is no reason for this to be the case in an industry where capacity can be freely added and withdrawn and the market price freely set.

Since the value of ETH is decoupled from GAS and therefore from the volume of transactions on the Ethereum protocol, an ETH bull could argue that ETH tokens could have an arbitrarily high value without compromising the cost-efficiency of operations on the chain. But let’s first agree that because of the GASPRICE mechanism¹⁴ the volume of transactions on the ETH blockchain and the scale of its adoption are not transitive to a high ETH token value. This point is important as observers often erroneously assume that a high volume of network transaction volume driven by all of the different potential uses of the Ethereum protocol will necessarily give the ETH token high value.

Let’s work through some numbers to see what in fact the utility value of ETH might be. Ethereum GDP (i.e., PQ) is the total ‘revenue’ of the computing network performing the

¹³ Today in practice it seems that the vast majority of transactions use the default 0.02 microETH price, but that most likely reflects the incipient nature of activity on the network. GASPRICE can be expected to become more market-driven as use of the Ethereum network grows. From a basic microeconomic perspective, if the GASPRICE (in fiat terms converted via the GASPRICE to ETH exchange rate and the fiat value of ETH) exceeds from time to time the actual fiat cost of providing the requisite computing resources, you would expect users to reduce GASPRICE offered or miners to add competing computing resources to the network until the marginal cost again equals the marginal revenue, driving a decline in the GASPRICE. This relationship should hold no matter what the scale of the operations being performed on the blockchain. The market will just keep allocating more computing and storage resources to the network as long as it is profitable to do so.

¹⁴ Note that the GASPRICE mechanism helps to reduce the incentives to fork the chain because economic rent can be eliminated quickly through it without necessitating a fork. Protocols without a GAS mechanism can be expected to end up at a similar economic equilibrium through forks as Ethereum will reach through the GAS mechanism. Ethereum may still fork for other reasons.

underlying operations, which can be directly measured as GAS used multiplied by the average GASPRICE. On 23 December 2017, the total amount of ETH used to fuel (pay miners for) transactions on the Ethereum network was ETH 1,388 (derived from the total GAS used¹⁵ multiplied by the average GASPRICE that day¹⁶)¹⁷. ETH 1,388 is worth about \$1 million at \$700 per ETH. Annualised (simplistically multiplying by 365), this is about \$355m per year.

We can then play with different assumptions for how fast the Ethereum network will grow vs the declining computing and energy costs. For example, let's assume Ethereum network traffic grows from here at the same rate internet traffic grew from 1995 to 2005 (roughly 150% growth per year)¹⁸ and that the combined offsetting impact of declining computing costs is -20% per year (optimistic as this approximates only the effects of the average rate of decline in computing costs without a change in the consensus mechanism; implementation of proof-of-stake or other scaling solutions could represent a step change down in the computing costs of the network by orders of magnitude). The combined net effect would imply 'Ethereum GDP' (PQ) doubles each year. At this rate, Ethereum GDP would grow from \$355 million to \$363 billion in ten years, an over thousand-fold increase. If we assume an ETH velocity of 7, the network value of ETH would be \$52 billion *in 10 years*, about 24% *less* than its current network value of approximately \$68 billion. Of course, in order to provide an attractive return to investors buying ETH today, its current network value would have to be significantly lower than \$52 billion (assuming investors would expect to make a 30 – 40% annual rate of return over that period, the current network value would need to be in the range of \$1.8 – 3.8 billion).

The foregoing calculation implicitly assumes that GASPRICE is already set at the level where miners are making zero economic rent and that Ethereum does not change its proof-of-work system, for example to proof-of-stake. As it's early days for Ethereum and mining computing resources are still catching up with demand, miners are probably still temporarily making positive economic rent, which means this back-of-the-envelope calculation in fact overstates PQ even if proof-of-work is maintained. More significantly, if Ethereum successfully moves to a proof-of-stake mining system and thereby substantially reduces the computational inefficiency inherent in proof-of-work where 99% of the computing power goes to proof-of-work and only a very small portion to actually maintaining the ledger, the PQ of the blockchain would fall massively and along with it the Ethereum network value. Recall also the analysis in the previous section explaining why staking of tokens for mining under proof-of-stake won't allow Ethereum to sustain a network monopoly premium.

Another way to look at this is to relate the Ethereum GDP to the total revenue of Amazon Web Services. AWS total revenue in 2017 is estimated to be \$16.8b, growing to \$40b in 2021 (according to JP Morgan), an order of magnitude smaller than our 10-year estimation for ETH GDP in the previous paragraph. If the velocity of ETH is 7, the Ethereum GDP (PQ of computational resources running the network) would need to reach approximately \$476 billion or 28 times AWS' current revenue to justify its current network value and excluding any return on investment during the years while Ethereum grows to reach that scale. Now, of course, AWS is just one provider of cloud services, but Ethereum is just one blockchain. Even if we assume that Ethereum will have some greater market share of blockchain than

¹⁵ <https://etherscan.io/chart/gasused>

¹⁶ <https://etherscan.io/chart/gasprice>

¹⁷ On 23 December 2017: GAS Used 41,686.74 million x Average GASPRICE 0.00000033285710975 ETH = 1,388 ETH.

¹⁸ <https://blogs.cisco.com/sp/the-history-and-future-of-internet-traffic>

AWS has of cloud, it is still hard to see how the current Ethereum network value can be remotely justified on this basis.

Note that in my reasoning about the future value of ETH at equilibrium, I have so far not taken into account the mined tokens that miners receive for performing computational services for the network as that value does not accrue to token holders. Rather the opposite. There are in fact two negative impacts of mining rewards on the value per token:

- Issuance of new tokens doesn't increase the total network value, just as printing fiat money doesn't make people collectively richer in real terms. The new issuance goes to the miners at a one-for-one cost of dilution spread across the value of all pre-existing tokens. This must also be true for the interest rate (BIR) paid on ETH tokens deposited in a proof-of-stake system. The new tokens generated to pay the interest dilute all existing tokens such that the effect on the overall total value of ETH tokens is neutral. Those engaged in mining will benefit from the interest earned while those not engaged in mining will suffer from the corresponding dilution. But the existence of this system does not drive growth in the network value of ETH and in fact drives devaluation of each ETH token at the rate of total interest paid in ETH divided by the total issuance of ETH.¹⁹
- There is a second, subtler negative effect of this new token issuance. It subsidises the cost of operating the network, which at competitive equilibrium puts downward pressure on GASPRICE, which in turn puts downward pressure on the value of ETH at a constant GAS <> ETH exchange rate.

The paradoxical combined effect is that the cost of new token issuance through mining rewards is effectively borne twice by non-miner token holders.

The implication of all of the foregoing is that, even if Ethereum is hugely successful, the value implied by its use as a backbone utility protocol is likely a small fraction of its current value. All of this raises a question for ETH bulls: why would ETH be arbitrarily valuable if it's not some scarcity in relation to the volume of transactions and operations on the chain?

One proposed reason has been that people will hoard ETH as a currency with which to make financial investments, for example in ERC20 token ICOs or DAOs built on the Ethereum protocol. In his blog post "Platform Currencies May Soon Be Obsolete"²⁰, Aleksandr Bulkin articulates why it is unlikely that a single blockchain will host a large number of Dapps and at the same time function as a major monetary store of value. Also, if utility protocols turn out to be poor financial investments as the foregoing analysis suggests, how much investment demand will there be? Finally, in a frictionless, multi-protocol future, why stockpile a particular token specifically to make a particular type of investment rather than store your value in the best pure store of value protocol (or in productive investment assets) and acquire the amount of ETH or any other currency for a particular purpose (including a subset of investment purposes) at the time of need?

So that leaves the possibility that ETH replaces Bitcoin and becomes the dominant non-sovereign monetary store of value simply on pure store-of-value merits. We'll go deeper into the topic of monetary store of value below, but from where we are today, an objective observer would give Bitcoin significantly higher odds than ETH of becoming such a store of value. And as for those who argue that you can recreate Bitcoin on top of Ethereum, the

¹⁹ Vitalik Buterin, [Incentives in Casper the Friendly Finality Gadget \(v 27 August 2017\)](#), p6.

²⁰ Aleksandr Bulkin, <https://blog.coinfund.io/platform-currencies-may-soon-be-obsolete-78d9b263d902>.

question is, why would you? Why substitute a new sub-token on top of a more complex protocol with a larger attack surface, shorter track record, less decentralised governance and propensity to make backwards incompatible protocol changes, for a hugely robust, stable, proven, and widely accepted protocol that already performs that narrow function very well?

Cryptoassets as Money

Money is a debt ledger with three sub-functions:

1. Store of value
2. Means of payment
3. Unit of account.

Cryptocurrency's performance advantage over incumbent forms of money is (a) strongest and most obvious as a monetary store of value; (b) stronger for some, but far from all, payments; and (c) differentiated as a unit of account for a few select purposes.

Cryptocurrency is overwhelmingly better as a monetary store of value than, say, gold. (I won't enumerate the reasons why, as it's pretty intuitive and has been written about widely.) As a means of payment, it can perform better than incumbent technologies in specific instances (think international payments), but Visa, Apple Pay, Google Pay, PayPal and fiat currency work well and better than cryptocurrency for most day-to-day payments. As a unit of account, a non-sovereign cryptocurrency could be most useful in international trade, global commodity markets, foreign reserves, and jurisdictions with unstable domestic currencies.

Before addressing the question of how to think about valuing the payment and the monetary store of value functions of a cryptocurrency, I'll first examine the link between payments and monetary store of value. Many observers presume this link to be very strong, but the reality is more nuanced.

First, let me draw a distinction between a monetary store of value and a run-of-the-mill asset. A monetary store of value is characterised by having a value that is decoupled from its utility for other purposes and from the cost of making/extracting and storing it. A warehouse full of goods, a stockpile of copper and a tankful of petrol are all assets and have value (determined by the market at the equilibrium point where their marginal utility meets their marginal cost of manufacture/extraction, i.e., $MR = MC$). Inventories of assets such as these appear on a company's balance sheet, but companies seek to minimise how much they have to hold to carry out their business, given the capital carrying cost. They don't try to accumulate these inventories to store their retained earnings. Gold, by contrast, is a monetary store of value. Its value is decoupled at equilibrium from the cost of extracting and storing it. While we may also use it for jewellery (an ancient way of signalling our wealth to other members of society), and we use a bit of it for manufacturing electronic goods and other industrial uses, we also store tonnes of it at great expense in giant inert lumps as a form of savings—a store of value—with no intent of ever using those lumps for any other purpose. Gold is therefore arbitrarily expensive relative to its extraction and storage cost. Its value is subjective.

Consider some examples of the things we use as means of payment versus those we use as monetary stores of value today²¹:

²¹ Note that I exclude here things like pre-paid debit cards, gift cards, pre-paid telephone plans and air miles as they are relatively immaterial to the financial system. As it happens, these can all be used for payments and are assets but people treat them as working capital (immobilised balance sheet assets with a carrying cost) rather than a form of savings, so if anything, they are more payment rails than monetary stores of value.

- Means of payment: Visa (credit and debit), SWIFT, PayPal, Apple Pay, Google Pay, Western Union, physical cash
- Monetary stores of value: Gold, fixed and demand bank deposits, physical cash.

What's interesting is that the only thing that appears as both a means of payment and a monetary store of value is physical cash. Yet even though physical cash is clearly both a means of payment and a monetary store of value, individuals typically hold only what's in their pockets and extract more from a deposit account as they need it. Companies that aren't retailers typically hold zero or close to zero actual physical cash (instead keeping their treasury in bank deposits, commercial paper, treasuries, etc.). For a retailer, cash in tills is not even treated as money but rather as working capital. Bank deposits are a monetary store of value but are neither a payment rail nor cash; rather they are contractual obligations of a financial institution operating on a fractional reserve model. When you make a payment, you can convert the deposit (store of value) into physical cash (payment rail) and pay with the cash or you use your Visa (payment rail), which is then paid by way of your local interbank payment network (payment rail) through a change in ledger entries of bank deposits held by you, Visa and the merchant. Credit cards clearly aren't a store of value. SWIFT is a payment rail but stores no value. On the other hand, gold is just a monetary store of value but not a payment rail. No companies keep their accounts and no retailers price their goods in ounces of gold. No one pays for coffee with gold. But that doesn't dissuade people from using gold as a store of value.

The reality is that means of payment and monetary stores of value are more generally separated than combined. The point of all of this is to illustrate that there are lots of means of payment which don't represent stores of value. It is thus overly simplistic to assume that people will hoard that which they use to make payments as opposed to converting their store of value via the payment rail at the time of payment in the exact amount needed and for as little time as possible.

There is substantial evidence that economic actors choose what to use as a means of payment and what to use as a monetary store of value somewhat independently of one another, based on the inherent functional merits and demerits of the 'thing' in question as a payment rail or monetary store of value. Cryptoassets are an interesting special case where, since the technology is potentially advantageous compared to incumbent forms of money as a monetary store of value and also for some payments, it is possible that a single cryptoasset might successfully compete as both. (The counterargument is that, because it is likely to be effectively frictionless to convert between two cryptocurrencies, it will be even easier to disaggregate the payment and store of value functions than it is with incumbent forms of money.)

A useful thought experiment in this regard is to imagine there are competing cryptocurrencies that have the following hypothetical relative utility scores across the five key characteristics of money:

<i>Utility (10 high - 1 low)</i>	Coin A	Coin B	Coin C
Scarcity	10	2	7
Durability	10	3	7
Portability	6	10	7
Divisibility	6	10	7
Acceptability	3	10	7
<i>Total</i>	35	35	35

When the dust settles, what will be the outcome? Which hypothetical coin will be the dominant payment rail and which one the dominant monetary store of value? Will Coin C be the one coin to rule them all?²² Or will Coin A emerge as the dominant monetary store of value and Coin B the dominant payment rail with users converting as required between the two? Much will depend on the technical and political trade-offs in creating cryptoassets that score relatively higher or lower on each of the five dimensions, but I lean towards the view that specialisation combined with protocol interoperability is a more likely equilibrium than one coin to rule them all.

It would be rash to go so far as to dismiss payments functionality as a necessary feature of a dominant monetary store of value (despite gold having none). A cryptoasset aspiring to be the dominant monetary store of value should prudently strive to have reasonably good payment functionality (divisibility, fungibility, acceptability); the more the better, provided that its monetary store of value functionality (scarcity and durability) isn't compromised. Poor or no payment functionality could impair a cryptoasset's ability to be adopted as a monetary store of value, so it's an important feature. But the point of all of this is to observe that a cryptoasset with the strongest monetary store of value functionality and with good but perhaps not the strongest payment functionality has a strong chance of winning out as the dominant store of value. Payments functionality in a monetary store of value cryptoasset is a satisficing, rather than a maximising, condition.

By corollary, just because a cryptoasset has an edge in payments doesn't mean it will automatically become a store of value. To wit, the only means of payment listed above that is also a store of value is physical cash. If a cryptoasset like Ripple is better for payments (cheaper transactions and more bank support) but weaker as a store of value than Bitcoin (due to centralisation of governance and supply uncertainty), it's unlikely that Ripple will win out as a store of value. And just because users employ some utility cryptoasset (such as ETH) for practical purposes, it doesn't mean that they will see it as a monetary store of value and hoard it as a form of savings rather than treat it as working capital to be minimised on their balance sheets by buying just as much ETH as needed when needed.

²² A crypto-spork? The crypto equivalent of that hybrid spoon and fork that no one uses because it's not as good as a fork in piercing solids nor as good as a spoon in transporting liquids. Insightful humour credit: <https://medium.com/@hamptonfischer/bitcoin-cash-a-spork-7f9f6230a57>. Or rather, a crypto-fpoon? But I digress.

What does all of this mean for assessing the potential value of cryptoassets as money? First, you should look at payment functionality valuation and monetary store-of-value valuation as two separate and additive things. You should then consider the relative functional strength of the cryptoasset you are valuing on payment and store-of-value dimensions compared to competing cryptoassets, to check that extreme weakness on one dimension doesn't result in weakness on the other dimension; you shouldn't just assume that payments and monetary store of value are inseparable. The network value of a hypothetical cryptoasset that ends up serving as both a monetary store of value and a payment rail can be thought of on a sum-of-the-parts basis, where (total network value) = (monetary store-of-value valuation) + (means of payment valuation).

What do I think could be the equilibrium outcome? A very credible scenario is that you end up with a mix of non-sovereign cryptocurrencies, sovereign digital currencies, off-chain/layer-two payment solutions and evolved versions of centralised payment systems such as Visa, PayPal and Apple/Google Pay, competing in payments (with each having different strengths and weaknesses for specific types of payments), along with a single dominant non-sovereign monetary store of value, playing a role much like that of gold today. That monetary store of value could also replace a significant portion of foreign reserves and perhaps become a unit of account for international trade and commodities.

It seems likely that payments will remain a fragmented market and that sovereign digital currencies, off-chain (censorship-resistant or not) payment solutions and centralised payment systems will successfully compete for the vast majority of global payment volumes, most of which are small, domestic and mundane, and where speed and cost matter more than strong censorship-resistance. In this scenario, non-sovereign, censorship-resistant, decentralised payment protocols could end up being a niche product for international payments, markets with a failed domestic sovereign currency or payments where censorship-resistance is important (capital controls, sanctions, political repression, illicit activity). I can imagine sovereign states creating regulations that favour use of sovereign digital currencies rather than non-sovereign cryptocurrencies for domestic payments to help them retain control over domestic monetary policy and taxation. It's worth highlighting that, while sovereign digital currencies will probably successfully compete with non-sovereign payment-focused cryptocurrencies for payment volumes, they will likely facilitate the emergence of a non-sovereign monetary store of value, since their existence will help to eliminate the on- and off-ramp banking issues and friction currently involved in exchanging fiat for crypto.

To preview what follows, monetary store-of-value functionality is potentially far more valuable than payments functionality. Thus, if a monetary store-of-value cryptocurrency also works well as a means of payment, that's just a little additional value upside. What matters far more to us as investors is store of value.

Payments

Proceeding from the general observations at the beginning of this paper, a means of payment, examined in isolation from the monetary store-of-value function, is just another utility protocol where the value of the token cannot decouple from $M = PQ/V$ of the computing resource maintaining the payment blockchain. Payments on a large scale must be cost efficient. To the extent a cryptoasset is both a store of value and major payment rail, in order to be cost-efficient and competitive as a means of payment, the payment part will have to be economically disassociated from the store-of-value function, such that the incremental M for payments still equals PQ/V for the computing resources facilitating the payment function and where V can go very high. This dissociation can be achieved by way of an explicit

mechanism like Ethereum's GAS; by way of various scaling solutions like off-chain transaction processing; and/or by emphasising transaction fees rather than block rewards in rewarding the mining network. The effect of something like Layer 2 transaction processing is to massively increase V (and reduce M) for the payment component of the sum-of-parts valuation of a cryptoasset.

Blockchain payments will in fact be worth (to token-holders) much less than their centralised counterparts like Visa, Apple/Google Pay and PayPal are worth today. It is incorrect to think that because a cryptoasset serves as a payment rail, owners of the token in the system would own something comparable to the enterprise value of, say, Visa divided by the number of tokens issued. Rather, at mature equilibrium, the network value of such a token would be $M = PQ/V$ where PQ is just the aggregate cost of the computing resources to run the chain (which may be thought of as the annual IT budget of an equivalent-volume incumbent payment system multiplied by some coefficient to adjust for the relative computing inefficiency of decentralised vs. centralised architectures) and V is of course some (probably high) velocity. The value implied by the correct valuation framework of $M = PQ/V$ is much, much lower than the enterprise value of the incumbents. Blockchain payments may disrupt and displace the incumbents to the enormous benefit of users, but the value of these protocols as expressed through their tokens will be much less than the value of the disrupted enterprises.

An additional factor to consider is the extent to which companies and individuals choose to keep an inventory of payment tokens as working capital on their balance sheets. The question is how strong this working capital driver will actually be. Companies and individuals hold very little physical cash. To the extent they hold cash-equivalents as a buffer against uncertainty, they hold it in the same currency in which they incur expenses, which will play to the advantage of sovereign digital currencies and incumbent payment systems unless retailers and suppliers begin to reprice directly in cryptocurrency units on a large scale. Except for users who make lots of international payments, who are primarily engaged in international or commodity-based trade or who operate in economies with weak domestic currencies, it may not make a lot of sense to maintain a significant stock of a non-sovereign-denominated payment rail cryptocurrency. In a crypto-native world, moving from one cryptoasset to another (for example from a store-of-value cryptoasset to a means-of-payment cryptoasset or converting between alternative payment cryptoassets) will be trivial, immediate and frictionless. What would the rationale be for holding an inventory of a given means-of-payment cryptoasset?

It's also worth reflecting on the difference between cash equivalents, such as deposits at a fractional reserve bank, and physical cash, and how that distinction reads over to a cryptoasset. If there is no fractional reserve banking system available for the cryptoasset, users may opt to store their value in yielding fiat-denominated deposits while keeping a low inventory of the payment rail cryptoasset rather than have a large holding of a cryptoasset that doesn't yield anything, or they may lend their cryptoassets out by buying yielding cryptoasset-denominated bonds and commercial paper.²³

Finally, in thinking about the potential value of the payments function of a single non-sovereign cryptocurrency, it's noteworthy that, while there clearly are network effects in payments (Visa is more useful and worth more as a business than Diners Club because it's accepted by more merchants and used by more consumers), payments today do appear to be a

²³ Note that both of these things increase the money multiplier and effectively decrease the scarcity of the cryptoasset in question.

structurally somewhat fragmented space. How many non-sovereign monetary stores of value are held in portfolios? Gold is pretty much it. Now think about how many different payment rails you've used over the past month: physical cash (perhaps in multiple currencies), Visa, Amex, PayPal, direct debit, SWIFT, etc. They were all good and reasonably fit for purpose in slightly different ways and with slightly different features for a specific payment: cash to tip the porter, Visa to pay on Amazon, Amex to buy a plane ticket and get the points, PayPal to pay on that dodgy website you don't trust with your card number, direct debit to pay your utilities bill, SWIFT for an international transfer. We should always be careful not to assume the new paradigm will function like the old paradigm (in the 90s, for example, we imagined old media online rather than social media). But we can arrive at the expectation that payments will be a fragmented space by applying first principles rather than extrapolating from the present. Moving among cryptoassets will be frictionless, and there are lots of nuanced differences across payment instances. It's thus more likely that we'll use lots of different payment rails based on how their respective features mesh with the circumstances. Time will tell what those will be, but it's not hard to come up with an initial list of possibilities: fully-autonomous smart contract payments that require a Turing-complete language overlay, payments where speed is of the essence or where cost is of the essence, payments where security or anonymity dominate, cases where one merchant simply accepts Litecoin while another takes Dash and many, many others.

In sum, I can imagine constant innovation and an ever-changing, fragmented, increasingly competitive payments landscape. This stands in contrast to monetary store of value, where leadership tends to strengthen over time. Path dependency is likely to be much stronger in store of value than means of payment. While there is some value in a payment rail exclusive of monetary store-of-value value, it is relatively low. Absence of payment functionality may hinder a cryptoasset's monetary store of value utility, but there are many examples where that isn't the case and where payment functionality doesn't translate into store of value, so any causality is weak.

Monetary Store of Value

I would argue that one cryptocurrency will likely become the dominant non-sovereign monetary store of value, because it's not clear what utility having two or more of them would add. Gold is the one dominant monetary store of value that isn't a fiat currency or tied to one. Yes, there's silver, but the value of all silver is a tiny fraction of the value of all gold, only about 20% of annual silver demand (worth about USD 3.3 billion in 2016) is for monetary uses²⁴, and it's trivial from a financial markets perspective. Why would we need multiple cryptocurrencies serving as non-fiat monetary stores of value? What utility would that add?

That brings us to the matter of quantifying the potential future value of the dominant store-of-value cryptoasset—and therefore the upside relative to today's value. If a cryptoasset becomes a dominant non-fiat monetary store of value, a logical place to start in estimating its potential network value is as a fraction or a multiple of the value of the total stock of the current technology filling that role, i.e., gold, which has a total value of USD 7.8 trillion.²⁵ The question of what fraction or multiple to apply is more subjective. You may feel it will be hard or take a very long time for a cryptoasset to fully replace gold, which has been around for millennia, in which case you think it will be a fraction. Or you may argue that the

²⁴ <https://www.silverinstitute.org/silver-supply-demand/>

²⁵ USD 7.8 trillion = total estimated gold above ground of 187,200 metric tonnes x 32,150.7 troy ounces per metric tonne x USD 1,292 per troy ounce on 21 August 2017, the date this section was written.

technical advantages in terms of divisibility and portability in particular will mean that more of the world's population will hold this cryptoasset than they do gold (anyone with a smartphone, a memory stick or a paper wallet can hold any quantity of a cryptoasset, but carrying and storing investment gold is more difficult). That cryptoasset, moreover, will likely play some role in payments while gold does not.

To try to bring some objectivity to this last question, let's look at the breakdown of where gold is today. Of total above ground stocks of gold of 187,200 metric tonnes, 38% is in the form of bullion holdings, of which a little less than half was held by the official sector (i.e., national treasuries) and the remainder by the private sector. The remainder of above-ground gold is almost all in the form of fabricated products, which breaks down roughly into 80% jewellery and 20% industrial products.²⁶

With those building blocks you can play with your own different scenarios of what success looks like. I'll develop a strawman scenario for consideration below.

While some jewellery may be notionally held for investment purposes, and there may be some collectible coins in the fabricated products number that holders think of as 'investments', I broadly exclude fabricated products, as a cryptoasset isn't a substitute for the vast majority if not all of those uses. Instead I focus on the bullion holdings. Because of cryptoassets' superior features over gold and its additional utility advantage for some payments, we might assume that a successfully dominant cryptoasset store of value being worth at maturity 1 – 3x private bullion holdings. Because national treasuries may be slower and reluctant to adapt, I'll suggest an assumption at a ten-year horizon of 0.25 – 1x official bullion holdings. Breaking the 38% of gold represented by bullion into 20% private and 18% official holdings and using the foregoing assumptions, we would estimate that the dominant store-of-value crypto could be worth 25 – 78%²⁷ of total gold stocks at maturity, i.e., USD 1.9 – 6.1 trillion.

Displacing gold bullion could, however, just be the tip of the iceberg. Gold represents a little less than 11% of the USD 12.7 trillion of total international reserves, with fiat currencies making up 86% and IMF Special Drawing Rights (SDRs) and IMF-related assets the remaining 3%.²⁸ The fiat currency portion is made up of 63% USD, 20% EUR and the remainder other currencies (most significantly GBP, JPY and CHF).

We need to separate fiat currencies' use for domestic payments from their use as international reserves. As stated above, there are good reasons to be sceptical about how significant a portion of domestic payments a non-sovereign cryptoasset will replace outside of countries with unstable sovereign currencies, namely: the existence of incumbent low cost and efficient centralised payment rails, the unwillingness of states to give up control over domestic monetary policy and the inevitability of sovereign digital currencies. But think how uncomfortable a situation it is for countries to hold the bulk of their international reserves in other countries' fiat currencies. We know that China and Russia in particular chafe at this situation. They would love to have an alternative to the USD and EUR and have even talked about creating an alternative (which didn't go very far, since creating such an alternative would require their trusting each other). Think about how difficult it must be for the Chinese to have exports often priced in USD and for commodity-producing nations that commodities are priced globally in USD. Furthermore, any USD-based transactions must pass through

²⁶ [GFMS Gold Survey 2017](#) p.36

²⁷ Low end: $((20\% * 1) + (18\% * 0.25)) = 25\%$; High end: $((20\% * 3) + (18\% * 1.0)) = 78\%$.

²⁸ As of 28 April 2017. [International Monetary Fund 2017 Annual Report](#). Note that on 28 April 2017 an [IMF Special Drawing Right = 1.371020 USD](#).

SWIFT, which is controlled by the US, and exclusion from SWIFT would be tantamount to near-complete isolation from the international financial system (witness Iran). This is an increasingly untenable situation for many sovereigns, particularly as the global power and influence of the US wanes. A non-sovereign, non-fiat, trustless, censorship-resistant cryptoasset would be a far better alternative for most foreign currency international reserves. IMF SDRs are already a synthetic store of value, so could also be easily and sensibly replaced by such a cryptoasset.

Building on our gold bullion analysis above to put some numbers around the potential implications for the network value of our monetary store-of-value cryptoasset, we might assume that it replaces somewhere between 0.25x and 0.75x of non-gold international reserves. My low-end assumption is fairly arbitrary but my high-end assumption reflects the likelihood that states will want to diversify their foreign reserves to some extent as they already do in holding fiat currencies other than the USD. These assumptions would add a further USD 2.8 – 8.5 trillion in value to our dominant monetary store-of-value cryptoasset. Adding these amounts to our gold bullion-based numbers above gives a total potential value range for our dominant monetary store-of-value cryptoasset of USD 4.7 – 14.6 trillion.²⁹

I'll stop there for now, but there are two further upsides that I haven't explicitly taken into account. First, it could make sense for such a cryptoasset to replace the USD as the standard unit of account for global trade and commodity prices. Trade- and commodity-centric firms may therefore choose to capitalise themselves in such a cryptoasset, creating further demand for its limited supply, mitigated by the inevitable emergence of fractional reserve banking and bond markets denominated in the cryptoasset which would increase its money multiplier. Second, such a cryptoasset will likely be used for some payments, such as international payments or domestic payments in countries without stable sovereign currencies (where this is already happening). This latter potential is at least somewhat, if not fully, captured in the high end of the range above in the sense that, when we start thinking about this cryptoasset store of value representing a multiple of private gold bullion holdings, we are implicitly already accounting for some displacement of physical cash holdings.³⁰ As explained in the previous section, the incremental sum-of-parts value contribution from the payments functionality arguably won't be that significant compared to the store-of-value component, so ignoring it here probably doesn't very materially impact the potential value target.

The next question is, which cryptocurrency has the highest probability today of becoming the dominant store of value? It seems to me that the probable answer based on the information in our possession today is Bitcoin (BTC). It has more users; has decentralised (to the point of dysfunctional) governance; has more hashing power than any other crypto; is highly stable and robust; has been around longest; and has never been hacked. Other cryptoassets may have features that Bitcoin doesn't have that are useful in sundry use cases other than store of value, but store of value is a simple functionality (perhaps the simplest of all the cryptoasset use cases), and Bitcoin has been and continues to acquit that functionality flawlessly. Critics point to the conflictual politics that complicate changes to Bitcoin's code, but seen purely through a monetary-store-of-value lens, that can be seen as more of a feature than a bug. It seems to me that it is far more likely that Bitcoin becomes the dominant store-of-value crypto than some other existing or future contender that isn't Bitcoin. If Bitcoin were to become the dominant monetary store of value cryptoasset, based on my total mature network value

²⁹ Total international reserves of USD 12.6 trillion, of which 89% non-gold reserves = USD 11.3 trillion. Low end: USD 11.3 trillion x 25% = USD 2.8 trillion; high end: USD 11.3 trillion x 75% = USD 8.5 trillion.

³⁰ To give some bounded idea of the potential value of displacement of some domestic fiat currency holdings, global M0 is about USD 5 trillion, so we'd be talking about a relatively small fraction of that.

estimate of USD 4.7 – 14.6 trillion, it would be worth approximately USD 260,000 – 800,000 per BTC fully-diluted³¹ at maturity.

We should pause here to think about how long the emergence of a cryptoasset as a dominant monetary store of value might take. On the one hand, gold has been around for millennia, so the mental paradigm shift required might take longer than 10 years and never occur fully. On the other hand, we rode horses for transportation for millennia and moved on from that pretty quickly and categorically with the advent of the superior technology of the motorcar. That transition required a major build-out of physical infrastructure while the one that interests us here requires little more than a shift in mindset. Also, financial markets tend to discount the future as soon as there is consensus about it, so the value per Bitcoin could anticipate the levels of adoption I'm holding up for consideration. (For a more extensive discussion of this topic, see below for **Addendum on Pace of BTC Price Rise**)

Regarding risk, an investment with a 20x – 60x upside³² only requires a probability of success of between 2% to 5% to be a positive net expected value investment. Each of us can reflect on his own view of what that the probability is of the foregoing scenario materialising. I'm personally pretty comfortable that, given where we are today in Bitcoin's development and adoption, that the probability is higher than 2 – 5%, likely much higher. While there are many technical, political, regulatory and psychological hurdles ahead, the store-of-value use case is by far the simplest one, and already closest to reality. I would argue therefore that here you have an investment with a downside:upside skew of -1x : 60x and a positive net expected value. Investments with both those characteristics are extremely rare.

While this paper isn't focused on analysing the risks ahead, it's interesting to observe that, of all the potential use cases for cryptoassets, monetary store of value is the one with the least technological risk. While Bitcoin will continue to evolve and improve over time (hopefully becoming more scalable, more fungible, etc.) and those improvements represent upsides, it doesn't in fact need to improve (or at least not materially) in order to replace gold and most foreign reserves. The existing state of the software and the existing network infrastructure is basically in place for this basic gold 2.0 / foreign reserves 2.0 function. Pretty much all that is required for that to happen is adoption and a change in popular and institutional perception and attitudes. In contrast, the more significant EVM-type ambitions of decentralised utility protocols require a number of technical advances and significant investments in infrastructure beyond what we have today. It is reasonable to believe those advances will occur in time, and we should all hope that they do, as they have the potential to make the world a better place, but it is obviously a longer, riskier path.

It is often proposed that Bitcoin's lead as the emergent dominant crypto monetary store of value could be usurped by another existing or future cryptoasset. This is true, but as Bayesians, we arrive at our views using probabilities based on the information at our disposal and update them as new information emerges. Based on the current information available to us, Bitcoin has the highest probability of becoming the dominant crypto monetary store of value, and that probability would appear to be high enough (greater than 5%) to make it a rational investment. As new information emerges regarding existing and new contenders for

³¹ To arrive at this target price, I have used a fully-diluted number of BTC of 21 million – 2.8 million = 18.2 million. While the total number of BTC that will ever be issued is fixed at 21 million, blockchain analytics firm Chainalysis [estimates](#) that 2.8 – 3.8 million BTC have probably already been permanently lost and are unrecoverable. Even if some of these BTC are later recovered, it is reasonable to expect that more BTC will be lost over time, so using the low-end estimate of 2.8 million for our purposes here seems reasonable.

³² At the time of this update in late December 2017, BTC was trading at around USD 13,000.

the monetary store of value crown, we can and should update our assessment. So far, as Bitcoin's price has risen, so have its odds of success. For now, Bitcoin appears to remain a rational bet.

The question of forks often comes up and whether forks of Bitcoin undermine its scarcity. Because they share same hash power, Bitcoin forks will either need to demonstrate some differentiated and valuable niche functionality compared to BTC, or they will wither and die in time. As long as BTC continues to perform well as a non-sovereign, monetary store of value, any such differentiated functionality will likely need to focus on less valuable use cases. It's possible that a Bitcoin fork finds such a non-store-of-value use case and survives with some relatively low, sustainable value reflecting the niche functionality it has addressed. It's also possible that the market assigns some fractional value for some period of time to an alternative store-of-value fork either irrationally (an ideological schism occurs and self-sustains for a little while) or as a kind of ace-in-the-hole back-up against some corruption of the main blockchain. In the end, the equilibrium outcome is more likely to be a single, dominant, monetary store of value, and it currently appears more likely that that will be BTC. For an investor who already owns BTC, the prudent investment strategy is simply to hold onto any forked versions that credibly appear to have a sustainable use case and value as they are received. For a new investor, it probably doesn't make sense allocating capital to prior forks due to their lower probability of success and their relatively niche potential value.

BTC v BCH

This is a good juncture to touch on the recent Bitcoin Cash (BCH) fork from Bitcoin (BTC) in August and the subsequent community ideological split behind them following the abandonment of the 2x fork in November. BTC appears to be focusing first on being a censorship-resistant store of value and improving its scalability over the long term, foremost through second-layer solutions; BCH is focused on immediate payments competitiveness through on-chain scaling. At the time of writing, BCH is cheaper for payments than BTC, acceptance of BCH for payments appears to be growing, and there are doubts in some quarters about both the timing and degree of success of BTC's second-layer scaling efforts. On the other hand, BCH is seen as weaker as a store of value and to have a weaker development team. BCH would just be another alt-coin as far as BTC is concerned but for (a) potential user confusion due to the similarity of the names and the ownership of the domain bitcoin.com by one of BCH's promoters, who actively claims that BCH is the 'true Bitcoin'; (b) the fact that BCH and BTC share and compete for the same hash power; and (c) the fact that BCH is being promoted by individuals with significant BTC holdings, who run large exchanges and wallet companies and hold sway over a significant amount (maybe more than 50% collectively) of hash power. Factor (b) raises the concern that, because BTC's difficulty adjustment is fortnightly while BCH's is daily, a significant increase in the price of BCH relative to BTC would cause hash power to swing away from BTC to BCH, slowing or halting BTC block times until the next BTC difficulty adjustment.³³ Factor (c) means that BCH's main backers may try to attack BTC by various means with a non-zero probability of success.

[This blog post](#) and [this blog post](#) go into scenarios in which the price of BCH could increase in the short run relative to the price of BTC. In a nutshell, if such price swings are temporary, so should be the disruption caused by them, and as a store of value BTC is less sensitive (store-of-value use implies generally larger, less time-sensitive transactions) to this kind of short term disruption than BCH is as a means of payment (generally smaller, time-sensitive

³³ You can follow the swings in hashrate between BTC and BCH [here](#).

transactions). Furthermore, this disruption in block times assumes transaction fees are held constant, but in reality transaction fees can adjust upwards to defend against short term attacks to keep hash power allocated to BTC and keep blocks moving, especially again because BTC functions more as a store of value than as a means of payment and as such is less sensitive to fluctuations in transaction fees.

The easiest and most prudent way to hedge against this risk is simply to own the same number of both tokens.³⁴ But if forced to get off the fence, the implications of the investment thesis laid out in this paper are that we should bet on store-of-value strength over means-of-payment strength, as we expect the former to be worth more than the latter over time. By focusing on the competitive and commoditised payments space, BCH is fighting for a place in a structurally fragmented use case where it doesn't really seem to do anything new or better compared to existing payment rail cryptocurrencies such as Dash and Litecoin,³⁵ not to mention the sovereign digital currencies that will soon appear; BTC is currently out ahead on its own as the leading crypto monetary store of value, a use case that is less prone to fragmentation and more likely to be dominated by a single cryptoasset. If BTC is the stronger store of value, it should remain more valuable than BCH and profit-driven miners should continue to allocate more hash power to it over the long run. Tortoise and hare-style, there's also a good chance that BTC's various second-layer development efforts will make it more relevant for payments, smart contracts and the like over time, providing potential upsides beyond the core store-of-value case.

Conclusion

Due to protocols being open-source, the ability to fork, the competitiveness of mining and the importance of relative cost to adoption levels, the value of utility protocol tokens will at equilibrium not decouple from an $M = PQ/V$ valuation, where PQ is the total cost of the computing resources required to maintain the blockchains. This value will likely be relatively low due to the very high potential values of V (velocity) and will be deflationary in line with deflation in the cost of processing power, storage and bandwidth and due to scalability-enhancing innovation.

Public blockchain technology is an incredibly powerful engine for creating significant user surplus, but that surplus will go to users, not to token holders or miners. Investing in utility tokens is in the end tantamount to investing to own a bit of the currency used to operate a big, commoditised, perfectly competitive SaaS business that itself earns no sustainable economic rent. There will be some value there, but perhaps not much. It is possible, and in fact quite reasonable, to construct very bullish protocol adoption scenarios where the equilibrium network values are very low and lower than the current network values of utility protocols such as Ethereum.

While the scale of use of utility protocols may be very large (caveated by the fact that the inherent redundancy of trustless, censorship-resistant consensus mechanisms compared to centralised ones makes them more expensive to operate and therefore only economically relevant for a subset of potential use cases), the potentially very high velocity suggests that this future mature equilibrium value may be thought of as something in the tens or hundreds of billions of USD in aggregate. A sizeable amount, no doubt, but perhaps not sufficiently

³⁴ The BCH:BTC price is currently hovering around 20% and even if BCH fails to win dominance it is likely to persist and retain value for some time, so the cost of the hedge isn't huge.

³⁵ Bitcoin Cash's main selling points relative to other payment rail cryptocurrencies seems to be that its name includes the word 'Bitcoin' and that holders of BTC on 1 August 2017 received BCH for free in the fork.

attractive compared to the current ~\$250 billion of network value of all alt coins combined to provide an attractive risk/return for investors from where we are today.

In the context of evaluating cryptoassets as money and in a world where value can be moved among protocols with little or no friction, a cryptoasset can be a monetary store of value without being most efficient for payments or a great means of payment without being a store of value. We can therefore look at the potential value of a cryptoasset's monetary store of value function separately from its payments functionality. Monetary store of value functionality will likely be one or two orders of magnitude more valuable than means of payment functionality.

Payments are likely to be fragmented and transaction volumes shared across a range of sovereign digital currencies, off-chain payment systems, centralised payment systems and multiple non-sovereign cryptocurrencies, each with their respective strengths and weaknesses for specific payment instances. This, combined with the fact that payment functionality is analogous to utility protocols and will therefore be valued on a $M = PQ/V$ basis, means that the means of payment value of any given cryptocurrency will be relatively low.

In contrast, the potential value of a winning monetary store of value protocol can be measured in relation to the total value of gold bullion and foreign reserves, suggesting a potential value in the USD 4.7 – 14.6 trillion range. If Bitcoin were to become that monetary store of value (and it currently appears to be the strongest contender by some margin), it could be worth USD 260,000 – 800,000 per BTC, i.e., 20 – 60x its current value. If one places a higher than ~5% chance of Bitcoin succeeding in this way, it is a rational and attractive investment for a long-term investor before considering other potential upsides stemming from payments and unit of account utility. Investing in other cryptoassets based on use cases other than monetary store of value appears less compelling.

Let's be clear. This could all go substantially to zero for various reasons. Being 'right' in an investment with a high risk of failure but a highly positively-skewed distribution of potential outcomes is about getting the *a priori* probabilities right (as adjusted for new information as it arises) and getting position sizing right. Provided you accept that Bitcoin's net expected value is positive, even marginally so, the right answer on position sizing isn't zero. Nor of course is it 100% of assets. For those stuck at the step of whether or not to invest, the logical thing to do is to move past that point and focus on position sizing. If you're more sceptical, invest less. If more confident, invest more. But even for the most sceptical, you might constructively ask yourself, why wouldn't you invest USD 1? Well, rationally, you probably would. Now how about USD 2? Repeat until you get to your Bayesian optimal position size. Given the significant risk of loss, in most circumstances the correct answer is probably a long-term, buy-and-hold, unlevered investment of a low single-digit percentage of assets (at cost).

Addendum: Thoughts on pace of BTC price rise

9 December 2017

How alarmed should we be about the recent rapid run-up in the price of Bitcoin? We read every day now that the speed of the increase itself is a tell-tale sign of a bubble.

Let's think further about how long it could or should take Bitcoin to reach its long-term equilibrium value and the shape of the path it might follow. When considering the potential or appropriate pace of price discovery for different kinds of cryptoassets, we again need to distinguish between utility protocols (including means-of-payment protocols) and a monetary store of value protocol.

When we value the shares of a growing company, we discount our expectations of future cash flows based on the expected growth path of the business in acquiring customers, building its team, building out its infrastructure for delivering value to customers, developing its technology or products, etc. at some discount rate that reflects our assessment of how much risk there is around the realisation of those expectations. This means that a company's shares can be overvalued at price x at time t_0 but undervalued at the same price x at time t_n . To frame this with an example, it is not necessarily inconsistent to say that Amazon was overvalued at \$107 per share in 1999 even though it trades at \$1,162 per share in 2017. Amazon has grown exponentially and navigated tremendous risks while competitors have progressively ceded market share and while the economy and the markets Amazon serves grew over those 18 years. If you had bought shares when they first peaked around \$107 in December 1999 and held them until December 2017 at \$1,162 per share, you would have made a 14% annual rate of return. In hindsight that would have been a very good investment, but there are a lot of other companies you could have invested in at 1999 prices that would have been bad investments and it was very difficult to know *a priori* that Amazon would be an exception. A 14% return arguably wasn't all that great, or certainly not excessive, based on a reasonable assessment of the risks in 1999. If you had sat down at the time to do a DCF valuation of Amazon in 2000 and were discounting a set of financial projections reflecting what actually happened, you would have reasonably applied a higher discount rate than 14%. You could have earned a similar return investing for example in a fairly pedestrian diversified leveraged buyout fund with arguably much less risk of loss of capital.

What does all this have to do with how long it could or should take a cryptoasset to appreciate in price? Similar to the shares in a company³⁶, the growth of the value of a utility protocol (including a payments protocol) should accompany the growth in the number of users, volume of use/transactions, buildout of the network (for example, merchants accepting a particular payment protocol or the installed base of IoT devices running a kind of smart contract), progress in following any development roadmap, etc. So, when we say ETH should be worth \$52 billion 10 years, we aren't saying it should be worth \$52 billion today. Lots has to happen to make it worth \$52 billion in 10 years. Accordingly, you should discount the \$52 billion back to the present using some discount rate reflecting the perceived risk of that actually happening, versus something better or worse.

A monetary store of value protocol works completely differently in terms of the potential (or even appropriate) timing and pace of price appreciation. Provided a monetary store of value protocol and the network running it is technically capable of acquitting its function as a monetary store of value (as arguably is already the case of BTC), the pace of it reaching its mature equilibrium value is as fast or slow as the pace of collective mindsets seeing it as such. This could take centuries or only as long as it takes synapses to fire. The value of a bar of gold or of a Picasso at a point in time is simply the value we collectively assign to it. Of course, that value might still evolve over time along with the accumulated wealth of our society and the size of our economy, but it is decoupled from some path of growing cash flows or expected cash flows.

If we relate this back to the potential equilibrium value of a dominant non-sovereign monetary store of value, there are two distinct steps, each of which could take a very long time or only days/weeks/months and with a long or short hiatus between the two. Of the overall potential value estimate of USD 4.7 – 14.6 trillion, USD 1.5 – 4.7 trillion is based on private sector holdings and USD 3.2 – 9.9 trillion on public sector holdings. The first step is

³⁶ The analogy to equity is used narrowly here in the context of price appreciation. As discussed above, utility protocols are akin to money supply, not equity, in other respects.

for private investors to reach a consensus that it is a strong monetary store of value. It could take a very long time for private investors to reach this consensus, *or* the network value could simply gap up to that level in a matter of weeks or months as adoption moves from just retail investors to institutional and retail investors. This is the stage we're in the middle of today with BTC. Until 2017, BTC ownership was predominantly a retail investor phenomenon, dominated by techie early adopters and some UHNWIs and family offices with connections to tech.³⁷ Then, over the course of 2017, we saw a proliferation of crypto-focused funds (still collectively only representing an estimated USD 2 – 3 billion across 119 funds³⁸). While both those groups continue to grow, we are now seeing the much larger collective firepower of mainstream hedge funds, family offices and (U)HNWIs starting to come in. This could very easily turn into a stampede for the entrance and value could very credibly gap up to at least the low-end private-sector target of USD 1.5 trillion in a matter of months, with subsequent growth to the high-end private-sector target of USD 4.7 trillion happening more slowly over the course of two or three years. The potentially rapid move to USD 1.5 trillion would imply a BTC price of USD 112,000 based on the current number of BTC³⁹ (potentially relevant given the short time frame being considered) or USD 86,000 fully-diluted⁴⁰.

The speed with which this move, especially the first leg of it, could happen is accentuated by the fact that Bitcoin ownership is concentrated; that most Bitcoin haven't changed hands since the price was in the double digits;⁴¹ and that these owners have high conviction and high long-term BTC price expectations and have already weathered tremendous volatility for years without blinking. Fewer than perhaps 1 million BTC effectively circulate at all, and any new money will be forced to compete mostly for those, so the propensity for price to gap up as new institutional money flows in is high.

The second step is adoption by public institutions as a store of value and as a replacement of gold and foreign fiat in countries' international reserves. Of course, this could take a very long time, even after the private sector has embraced and accepted it. Governments move slowly. Decision processes are political. On the other hand, as soon as one government is known to have bought its first Bitcoin into its reserves, we could see a second stampede for the entrance as national treasuries around the world realise they need to diversify at least some of their reserves into Bitcoin before all their rivals do or be left at a strategic disadvantage. So here again, we could see a rapid addition to total network value of the low-end public-sector network value estimate of another USD 3.2 trillion, with that subsequently growing to the high-end target of USD 9.9 trillion (in addition to the private sector value above) over a few years.

The above scenario suggests that the price of BTC could increase by a very steep slope in the very near term as network value moves quickly to USD 1.5 trillion, then a flatter slope as private sector adoption matures towards USD 4.7 trillion and before the first adoption by countries as a component of international reserves. Then we could see another steep slope addition of USD 3.2 trillion to BTC's network value as countries stampede to the entrance, followed by a slower slope appreciation to USD 9.9 trillion of public sector adoption.

An interesting aside is that, with the shares of a company or the price of a utility cryptoasset, the higher the price from time to time, the higher the risk. For Bitcoin, paradoxically, the

³⁷ And people lucky enough to count Wences Casares as a friend.

³⁸ <https://www.hfalert.com/search.pl?ARTICLE=175427>

³⁹ 16.7 million BTC in existence less estimated 2.8 million lost BTC = 13.9 million BTC.

⁴⁰ 21 million BTC maximum less estimated 2.8 million lost BTC = 18.2 million BTC.

⁴¹ <https://bitinfocharts.com/top-100-richest-bitcoin-addresses.html>

opposite may be true up to a point and for certain periods. Right now, Bitcoin is still an insignificant curiosity to the financial markets, with a network value of just USD 218 billion. If and when it breaks through the USD 1 trillion level, it will likely be seen and accepted as a fully-fledged asset class and the financial world will go about building out complete financial markets infrastructure (full suite of derivatives, more robust and liquid exchanges and trading platforms, more custody options). That, combined with broadening private and public institutional ownership over time, will serve to consolidate Bitcoin's position as a monetary store of value and reduce risk. Of course, if at some point price starts to overshoot the potential value estimates for the relevant stage of adoption we're in, we could reasonably begin to worry. But let's cross that bridge if and when we come to it.

Consistent with the overall approach of this paper, the point of the foregoing isn't to say that the above necessarily will happen but rather to look at what reasonably could happen. Depending on what stage of the adoption path we are on and where we are relative to long-term potential network value and the interim milestone value points along the way, it may be entirely reasonable for the price to move up by a lot in a very short period. Such moves don't necessarily reflect irrational, bubble behaviour.