Watermarked tokens and pseudonymity on public blockchains

By Tim Swanson

Published: November 4, 2015

Abstract

Over its nearly seven-year incubation, Bitcoin has attracted a diverse set of entrepreneurs, developers, investors and researchers. While the initial attraction revolved around the recreation of a pseudonymous bearer asset, over the past year there has been a steadfast attempt to enable the Bitcoin network to transfer and secure off-chain registered assets. To achieve the goal of issuing, tracking and transferring virtual representations of off-chain assets, several startups have begun to provide watermarking services, commonly referred to as "Blockchain 2.0." This paper will show that a public, distributed ledger (such as Bitcoin) that secures off-chain assets cannot be both censorship-resistant and legally authoritative.

Consequently, a ledger that does not provide a one-to-one correspondence between what the endogenous network says and what the exogenous jurisprudence says about the status of a financial contract is a network that cannot exist without the legacy settlement framework that it seeks to replace, for the latter will continue remain the authoritative record of ownership. In practice, the censorship-resistant aspect of "Blockchain 2.0" is impractical as a solution for financial settlements in cash, securities and other off-chain property titles.[1]

Table of Contents

Abstract 1

Section 1: Background on the term *watermarked tokens* 3

Section 2: Registered assets conflict with Bitcoin's security assumptions 7

Section 2.1: How ungated overhead may lead to a "Top Heavy" scenario 11

Section 2.2: Turning it to eleven 17

Section 2.3: What is merged mining and why it is considered important 19

Section 3: Traceability and permissionlessness 23

Section 3.1: Government sanctions 27

Section 3.2: Encumbered titles 34

Section 3.3: Uniform Commercial Code 36

Section 3.4: FinCEN impact on watermarked tokens 38

Section 3.5: The Securities and Exchange Commission 40

Section 4: Conclusion 42

Acknowledgements 44

Disclaimer 45

Appendix A: Internal governance 45

Appendix B: A straw man game theory model 54

Endnotes 60

Section 1: Background on the term *watermarked tokens*

This paper is meant to provide readers a general overview of what watermarked tokens are, and what challenges exist for financial institutions looking to utilize these tokens on public blockchains. Specifically, it presents the argument that it is shortsighted and ill advised to try to secure off-chain assets with Bitcoin or other public cryptocurrency-based platforms.

There are at least three identifiable reasons that financial organizations looking to use some kind of public blockchain should be wary of a watermarked approach:

1) the built-in security system inherited from Bitcoin and other proof-of-work-based blockchains is not exportable in a regulated financial settlement setting (through a distortion of incentives);[2]
2) the lack of legal settlement finality; and
3) the regulatory risks that a watermarked approach introduces (see details in section 3)

This paper prefers to use the term **watermarked** token to encompass two types of systems:

1) Colored coins and
2) Embedded consensus systems which use their own proprietary metacoin

The current industry typically conflates colored coins for metacoins (and vice-versa) but they are technically not the same. Therefore, the term "watermarked" could be seen as a more inclusive umbrella term at this time.

Over the past two years, considerable marketing and advertising has attempted to highlight new functionality in the form of 'watermarking' which promises to expand the extensibility of public blockchains such as Bitcoin. The end goal is to enable Bitcoin (the network) to go beyond tracking ledger entries of just one specific on-chain asset (e.g., a bitcoin) to also allow users to watermark a bitcoin (sometimes referred to as a "token") to purportedly represent a sundry of off-chain assets.[3]

While the term "token" has a number of distinct definitions, for the purposes of this paper a **token** is limited to a discrete amount of unspent transaction outputs (UTXOs) that is tracked on a distributed ledger (which in this case is a blockchain).[4] These watermarked tokens act as representations for assets (e.g., common stock, gold, car deeds) that are exogenous to the ledger itself. Or in short, a watermarked token in this manuscript is a discrete amount of bitcoin that is tracked on the Bitcoin blockchain and is supposed to represent external value.

For instance, according to the first whitepaper on colored coins by (Rosenfeld 2012), these watermarked bitcoins can purportedly represent "alternative currencies, commodity certificates, smart property, and other financial instruments such as stocks and bonds."[5] Some literature (Schroeder 2015) refers to these types of instruments collectively as "cryptosecurities."[6]

Colored coins

There are no fewer than four startups that have attempted to capitalize upon the colored coins form of watermarking, including ChromaWay, CoinPrism, Colu and CoinSciences.[7][8]

Over the past year, each company has diverged and come up with their own proprietary standards, yet at the core the original idea is roughly the same: watermarking or applying "graffiti" to a specific amount of outputs – bitcoin outputs – typically measured in relatively *de minimis* quantities of satoshi (e.g., 0.0001 BTC).[9]

These colored outputs then conform to a set of rules, called a "color kernel," which associates certain arbitrary properties to these seemingly homogeneous units.[10] The bitcoins themselves are solely used as a carrier vehicle for other external value.[11]

The aim of these projects is to enable users to take a fraction of a bitcoin and "color" it to represent, for example, the title to a 2010 Toyota Camry, or shares of Twitter Inc. In theory, the private key holder can then transfer that asset along a public blockchain (such as the Bitcoin network) directly to other individuals or companies.

Based on marketing material, instead of having to transfer tens, hundreds or thousands of bitcoins in exchange for a good or service, users can instead exchange and manage entire asset classes in a relatively decentralized framework with a minute amount of colored bitcoin.

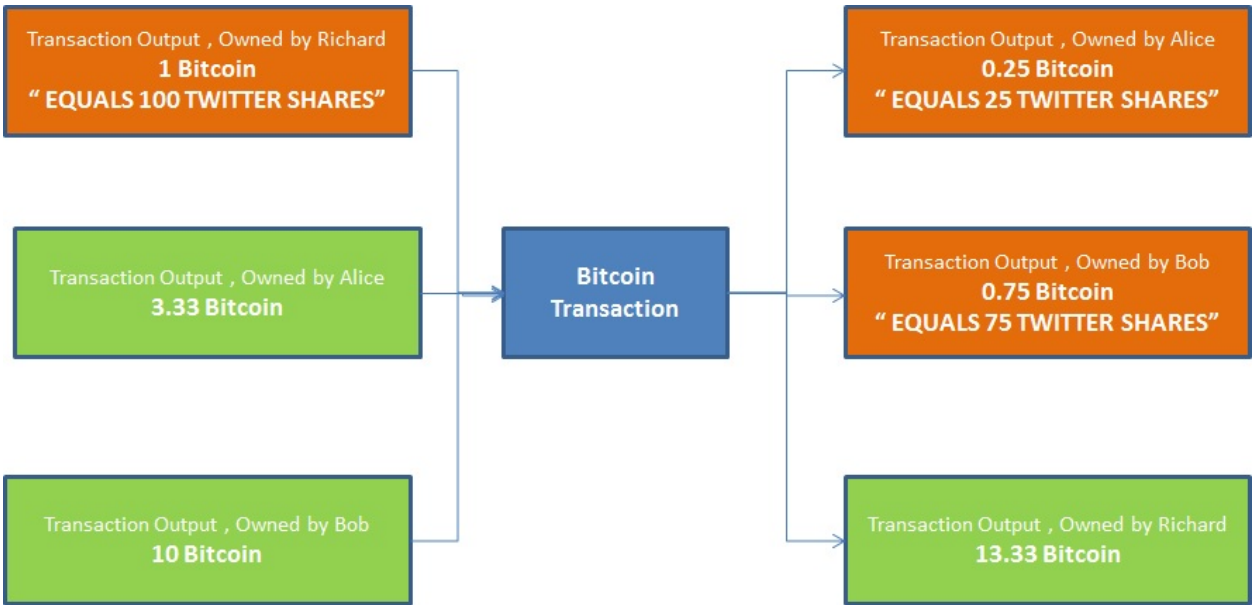

Simple model of a "colored coin." Source: Richard Brown[12]

It remains uncertain what the legal effect, if any, would be of Bob trying to transfer title of an off-chain titled asset vis-a-vis a public blockchain. At the time of publication, it does not appear that any Departments of Motor Vehicle or other official property system currently recognizes

watermarked, virtual tokens as proof of ownership let alone those with "bearer" characteristics.[13,14]

A **bearer**, in this manuscript, refers to an environment in which, in order to control an on-chain asset (in this case a bitcoin or colored coin) all Bob has to do is control the private key that corresponds to a specific amount of bitcoin. Because the Bitcoin network itself has no internal mechanism to tell if any other internal or external party has a better claim on the key, possession or knowledge of a private key amounts to *de facto* on-chain ownership.

This aspect of Bitcoin's structure (lack of gating and mapping of on-chain identities to off-chain identities) was purposefully done to enable commercial interactions between pseudonymous and anonymous parties. As a consequence, Bob does not have to prove his identity to the network in order to use the key. The party (or bearer) who controls the key can use the key and possess anything on-chain that the key directly controls.

This lack of recognizable identity between on-chain and off-chain environments represents a hurdle for financial institutions.

For instance, with respect to the transfer of ownership of common stock, (Schroeder 2015) points out a legal issue:

> Identification of ownership is perhaps even more important for blockchain stock because the rights to receive notices, to vote, to receive dividends, exercise appraisal rights, etc. are limited to registered owners, so there must be some way of identifying security owners. [15]

The intersection with jurisprudence will be looked at further in Section 3.

In practice, most coloring systems no longer watermark a specific set of satoshi, but instead use the OP_RETURN scripting code to encode or embed information onto a blockchain (usually the Bitcoin blockchain).[16]

Metacoins

The second form of watermarked tokens are embedded consensus systems (ECS), or overlay networks, such as Counterparty and OmniLayer (formerly known as Mastercoin) all of which embed themselves into a host network.[17, 18,19,20] Unlike colored coins, which are associated with specific unspent transaction outputs (e.g., a 'colored' satoshi), ECS are embedded messages associated with addresses.[21]

While they have different developer communities, these ECS each have created a new platform for asset issuance. Yet they simultaneously rely on the same set of host validators (miners) to secure multiple systems: Bitcoin mining pools secure not only Bitcoin's blockchain but also the Counterparty system and OmniLayer. In addition, ECS have issued their own set of native

cryptocurrencies, dubbed “metacoins” (XCP and OMNI respectively), which effectively corresponds with a small amount of satoshis on the host network.

What is the motivation for doing this? Why have these coloring and ECS projects been created?

In (Mizrahi 2015), the author notes one alleged advantage of using colored coins via the Bitcoin blockchain:

> Reduced implementation cost: we only need a thin layer on top of Bitcoin, there is no need to implement a cryptocurrency from scratch. [22]

Colored coins, metacoins and other graffiticoins were originally proposed as a “hack” that could rapidly bootstrap by using the Bitcoin blockchain for consensus without building an entire network from the ground up.

Since Bitcoin does not natively support the type of asset issuance or financial contracts that are required for financial applications, and because Bitcoin Core is not quickly amendable to such requirements (due to coordination problems discussed in Appendix A), external groups looking to quickly ramp up and use a public blockchain to create an off-chain asset tracking system must implement their protocol on top of the existing network.

Layers

As described above, colored coin issuers watermark messages by moving small amounts of bitcoins between addresses.[23] This solution is not particularly elegant or scalable, since the host network was not originally designed for this. In addition, it inherits all the information security safety issues associated with the host blockchain itself.[24]

The rigorous, more scalable, approach to creating an asset-tracking system using a blockchain is to create a robust protocol that includes asset issuance and native tracking.[25] The challenge is that this requires more programming effort, funding and skilled talent – all of which is scarce and has been difficult to procure by most of the watermarking projects.[26]

Example strawman model of a watermarked blockchain
Layer 1: Network
Layer 2: Consensus
Layer 3: Transaction
Quasi layer 4: Watermarked token

In short: ideally all asset issuance and tracking on a public blockchain should be native and endogenous to layer 3 instead of a jittery layer 4.[27] This is discussed in section 2.

What about direct forks of Bitcoin intended for private or permissioned use?

Though colored coins may work as a weekend project to understand how a blockchain and asset issuance works, the argument for colored coins makes little sense in a gated or permissioned system. If the fork is being brought in-house by an organization, the new developers can simply modify all of the code to natively include appropriate attributes, as this approach bypasses the coordination problem currently facing the Bitcoin Core community.[28][29]

While there is a temptation to take the shortcut to simply copy and modify the existing Bitcoin codebase, out-of-the-box, it does not immediately address the use-cases of financial institutions. For example, from the ground-up, some of the design principles of public blockchains such as Bitcoin are inconsistent for an architecture that is suitable for regulated financial institutions. The more plausible reason for Bitcoin forks is business related: Alice's startup does not have the resources or technical expertise to roll out their own new product that would integrate asset issuance and tracking in layer 3.

Building a new network with native issuance design parameters is no trivial feat. For instance, Ethereum raised (at the time) roughly $18 million and has a small army of developers (both paid and volunteer) to build a new type of blockchain from scratch and it still took them over one year to build a system that learned from Bitcoin's limitations in layer 3.[30] At the time of writing, few of the "private Bitcoin fork" projects being proposed have had access to similar scale of resources and have not made similar, materially significant code rewrites and changes to Bitcoin Core in their product.[31]

Section 2: Registered assets conflict with Bitcoin's security assumptions

The following section will look at the distortions caused by adding exogenous value onto a network that cannot detect and dynamically protect the exogenous value.[32] It also looks at what has empirically developed over the past year.[33]

On May 11, 2015, Nasdaq announced that it was experimenting with Open Assets, a colored coin implementation from CoinPrism, to record digests (hashes of summaries) of its private, pre-IPO Private Market™ book onto the Bitcoin blockchain – basically cap tables on a blockchain.[34] A month later it was announced that it is doing so in partnership with Chain.com, which had recently pivoted from the API space and into financial services.[35]

In October 2015 both Nasdaq and Chain subsequently announced that they are no longer using colored coins and are now using a new, non-Bitcoin blockchain for their projects.[36]

Prior to that announcement, one solution that Nasdaq, Chain and others such as tØ (tee zero) and Digital Asset Holdings may have looked at is using a hybrid multisig operation.[37] As described in (Mizrahi 2015):[38]

> A hybrid approach is also possible: a colored coin which represents property ownership will send to 2-of-2 multisig address, which requires signatures both from the registry and from the owner to unlock. In this case owner cannot transfer his property without interaction with registry, however, neither can registry do transactions without owner's consent. This scheme can provide extra security: a registry can perform additional authentication steps to make sure that transfer is correctly authorized.

But using a public proof-of-work ledger in such a centralized manner seems counterproductive.[39,40] To fully understand why, it is worth looking at some of the design assumptions.

For instance, (Sams 2015c) articulates several potential issues with the hypothetical 2-of-2 multisig model:[41]

> [A]dvocates of putting property titles on the bitcoin blockchain will object at this point. They will say that through meta protocols and multi-key signatures, third party authentication of transaction parties can be built-in, and we can create a *registered asset* system on top of bitcoin. This is true. But what's the point of doing it that way? In one fell swoop a setup like that completely nullifies the censorship resistance offered by the bitcoin protocol, which is the whole raison d'etre of proof-of-work in the first place! These designs create a centralised transaction censoring system that imports the enormous *costs* of a decentralised one built for censorship-resistance, the worst of both worlds.
>
> If you are prepared to use trusted third parties for authentication of the *counterparts to a transaction*, I can see no compelling reason for not also requiring identity authentication of the *transaction validators* as well. By doing that, you can ditch the gross inefficiencies of proof-of-work and use a consensus algorithm of the one-node-one-vote variety instead that is not only thousands of times more efficient, but also places a governance structure over the validators that is far more resistant to attackers than proof-of-work can ever be.

Recall that one of the design assumptions underlying Bitcoin is that the transaction validators – miners – are unknown and cannot be trusted, plus they may come and go over time (i.e., in theory there would be no fixed set of validators).[42] In fact, it was assumed that, because it is a public network, at any given time there would be different types of "mutually-distrustful peers" and all parties existed in the "presence of an actively malicious adversary."[43]

Similarly, all interaction and participation was assumed to take place – by default – pseudonymously. If you do not know any participants, let alone trustworthy participants, is there a way to secure this type of open network?

Bitcoin's solution to this question is a cost structure custom-built for maintaining a public, untrusted network and comes in the form of high marginal costs.[44]

This was put in place since anyone can submit transactions, even resubmit them multiple times (e.g, "double spending"). Therefore the network would be continuously open to spam attacks dubbed Sybil attacks or "pseudospoofing."[45,46] Due to all the rewriting, it would lack a canonical or definitive record of "truth."

Bitcoin sidesteps (but does not really "solve") the double-spend problem by purposefully making it expensive for any one actor to unilaterally change history (also called "the *state* of the ledger").[47]

Or as (Garay, et al. 2015) explain, "Nakamoto's protocol does not quite solve [Byzantine Agreement] since it does not satisfy Validity with overwhelming probability."[48] Validity is defined as the ability to "capture security/correctness against adaptive adversaries."[49]

To dissuade Sybil attacks, the Bitcoin network requires participating validators to submit proofs-of-work, a process labeled as "grinding," to prove that they have incurred a heavy cost and are therefore signaling to everybody else that they have paid a 'forfeitable bond' if others decide their work is invalid.[50] In other words, Bitcoin validation was *designed to be inefficient* in order to make attacking it economically expensive.[51,52]

Consequently, despite claims at Bitcoin industry conferences, there is no secure method of making proof-of-work-based mining efficient, or cheaper; for if the costs of maintaining the network were to decline, so too would the costs to successfully attack it.[53]

For better and for worse, Bitcoin and Bitcoin-like systems must be energy intensive, otherwise attackers could easily rewrite history. Miners compete through wealth destruction, as "real economic goods (time in fabs, electricity, engineering efforts) are being removed from the economy for the sake of proof-of-work mining."[54]

Or to quote (Cohen 2015):[55]

> Bitcoin involves proofs-of-work. There is no such thing as an efficient proof-of-work. That's a contradiction in terms. Bitcoin is designed fundamentally so that if people become more efficient at doing the proofs-of-work, the difficulty of proofs-of-work goes up right in tandem with what they're doing. And it ratchets so that your limit it will always be power.

At the time of this writing, the aggregate costs to maintain the Bitcoin network is around $300 million per annum (based on current token value). The largest component is the operating cost via electricity, which on the aggregate consumes several hundred megawatts of energy per year.[56] If market prices declined, less energy would be used but the network would also be *less* secure under the existing proof-of-work model.[57]

What does this design assumption, and the costs therein, have to do with securing off-chain titles?

By requiring identification and permission – by removing pseudonymity and turning a bearer asset into a registered asset as explored by (Mizrahi 2015, Bitfury 2015b, Bitfury 2015c) – Nasdaq and others are nullifying the primary utility of Bitcoin: censorship-resistance.[58] Yet the high marginal costs to maintain the underlying public network still remain.[59]

By grafting a permissioned setup onto a nominally decentralized, public blockchain such as Bitcoin via mandatory multisig and identification, the system not only becomes centralized but also forgoes the core benefit from a $300 million censorship-resistant network.

What does feature nullification mean?

(Schroeder 2015) points to this feature nullification as well. For instance, if Bitcoin holders in the US attempted to use a securities intermediary to satisfy the requirements of Uniform Commercial Code (UCC) Article 9:

> [I]t would defeat one of the primary advantages of cryptocurrencies over conventional payment systems such as checking accounts, credit or debit cards, automatic clearing house transactions, wire transfers or money transmissions. This is precisely the ability to engage in person-to-person transfers of value [without] the need to use a mediating bank, broker or other institution.[60]

Intersections with UCC are discussed later in Section 3.3.

On the face of it, trading and attempting to secure high-valued transactions is marketed as an acceptable use-case for reusing an existing public blockchain. However, the setup proposed above – permissioned-on-permissionless – is Rube Goldbergesque. Why not just use a permissioned ledger directly or even a database?

From a historical perspective, the idea of electronically managing a private cap table in near real-time is not new.[61] Founded in 2012, Mountain View-based eSharesInc has raised $25.8 million enrolled over 1,000 venture-backed companies by providing a similar service, albeit without the blockchain marketing moniker.[62] Similarly, prior to closing in January, Capography also attempted to provide a cloud-based cap table management service for startups.[63]

One last consideration regarding the Mizrahi model: an attacker on the Bitcoin network can still reverse an interval of "settled" transactions. The attacker's objective is undermining confidence in the market in which he has a short position (e.g., using Gemini), not double-spending a metacoin (which of course he cannot get away with because the identities of all counterparts are known to Nasdaq).[64,65]

As explained in (Sams 2015d):

> [T]his benign state of affairs is unlikely to persist if the bucket shops, which are today the only avenue for shorting bitcoins, are eventually replaced by professional derivatives markets. And it will certainly go away if billions of dollars worth of securities are represented through meta protocols on the bitcoin blockchain as some have eagerly extrapolated from the Nasdaq announcement. For then, attackers will have a way of constructing a scalable payoff for attacking the network: shorting the market in size. Acquiring a substantial portion of the network's hashing power is not an insurmountable goal. What is required is a sufficiently large monetary incentive to execute the attack. Putting billions of dollars worth of financial assets on the bitcoin blockchain materially changes an attacker's incentives.[66]

It bears repeating that both Nasdaq and Chain have announced that they are no longer using colored coins or the Bitcoin blockchain. There are a handful of others that still are though.

Section 2.1: How ungated overhead may lead to a "Top Heavy" scenario

Proof-of-work-based blockchains such as Bitcoin are, in practice, secured by proportionalism.[67] In the long run, it takes a bitcoin to make a bitcoin. This means that, with a few exceptions, rational laborers (miners) will not spend more than a bitcoin to make one. Or, in economic terms, the marginal cost (MC) of mining a bitcoin eventually equals the marginal revenue (MR) of its market value.[68]

The metacoins and colored coin projects listed above unquestionably increase the social value of the chain, yet they do not proportionally incentivize security beyond the existing block reward (seigniorage) subsidy. This could lead to an economic incentive to attack the chain, a type of fat tail risk that could dramatically impact any layer residing on top of the Bitcoin network.[69]

Is this an economic flaw? Weren't the mining mechanics of the Bitcoin network an incentive compatible Nash equilibrium, such that deviating from the assumed behavior did not result in a net gain for other parties?[70] Are there other incentives that change mining strategies to the benefit of those participants that "misbehave?"[71]

As mining rewards were established with the genesis block in 2009, Bitcoin and its progeny provide a fixed income on a regular timetable to miners. Consequently, using simple calculations, miners can gauge the potential short-term viability of their mining activities with respect to their hashrate relative to the market value of the token.[72,73,74]

Thus, if a bitcoin is worth $300, based on the current block reward, miners as an aggregate will not spend more than $45,000 per hour to secure the Bitcoin blockchain from reorgs and double-spending attempts.[75]

However, platforms that utilize and sit on top of Bitcoin's blockchain provide disproportional rewards because mining pools are effectively "color blind."[76]

In (Rosenfeld 2012) the author recognizes this as a shortcoming, noting that:[77]

> […] the token are any bitcoins that can be traced back to a particular output, and the transactions in which tokens are moved are Bitcoin transactions, recognized as normal transactions by oblivious Bitcoin nodes but must satisfy additional requirements to be considered legitimate by color-aware nodes.

Similarly, (Mizrahi 2015) states that some purported advantages of the Bitcoin blockchain are:[78]

> Security: By piggybacking on top of Bitcoin, we inherit its security properties. It is crucial for such a system to be tamper-proof, and Bitcoin is currently the cryptocurrency which is secured by the largest hashrate.
>
> Persistence: Significant total value of all bitcoins in existence create a huge financial incentive to keep the system alive and healthy.

On the face of it, colored coin systems and ECS such as Counterparty and OmniLayer require participants to pay some *nominal* transaction fee. Drilling down deeper, both are purposefully piggy backing and free riding off the block reward (seigniorage) provided to the mining network.[79] And currently, there is no native mechanism that requires or forces the fees to converge towards funding the gap between exogenous social value and endogenous on-chain value.[80]

Furthermore, as the historical record has shown with lifecycle of altcoins, or alternative cryptocurrencies, it cannot be assumed that there will always be a "huge financial incentive" to keep the system alive and healthy.[81]

(Buterin 2014) highlights this potential type of tragedy of the commons.[82] This is problematic for watermarked token issuers and users.

But can't it quickly be solved with higher mandatory fees?

(Rosenfeld 2012) mistakenly predicts a future empirical event:[83]

> By having the assets embedded in the Bitcoin blockchain, the entire existing infrastructure of Bitcoin network nodes and miners can be utilized to benefit each and every one of the assets. This allows creating a secure asset with no barrier to entry. The extra burden on the host network can be paid for with a property transaction fee system. The inclusion of arbitrary assets in the Bitcoin network will increase its total economic value, attracting more nodes and miners to strengthen it even further. Payments to asset holders can be done directly to the holding address in either the host or any guest currency.

The colored coin white paper above was originally published in December 2012 and many developments have occurred since then.

Yet two things that have *not* occurred are:

1) the creation of a "property transaction fee system" and

2) the subsequent attraction of "more nodes and miners"

In contrast, fees to miners are still an unsettled topic engendering enormous controversy from all corners of the ecosystem. Empirically fees denominated in fiat (such as USD), are roughly at the same level as they were in mid-2013.

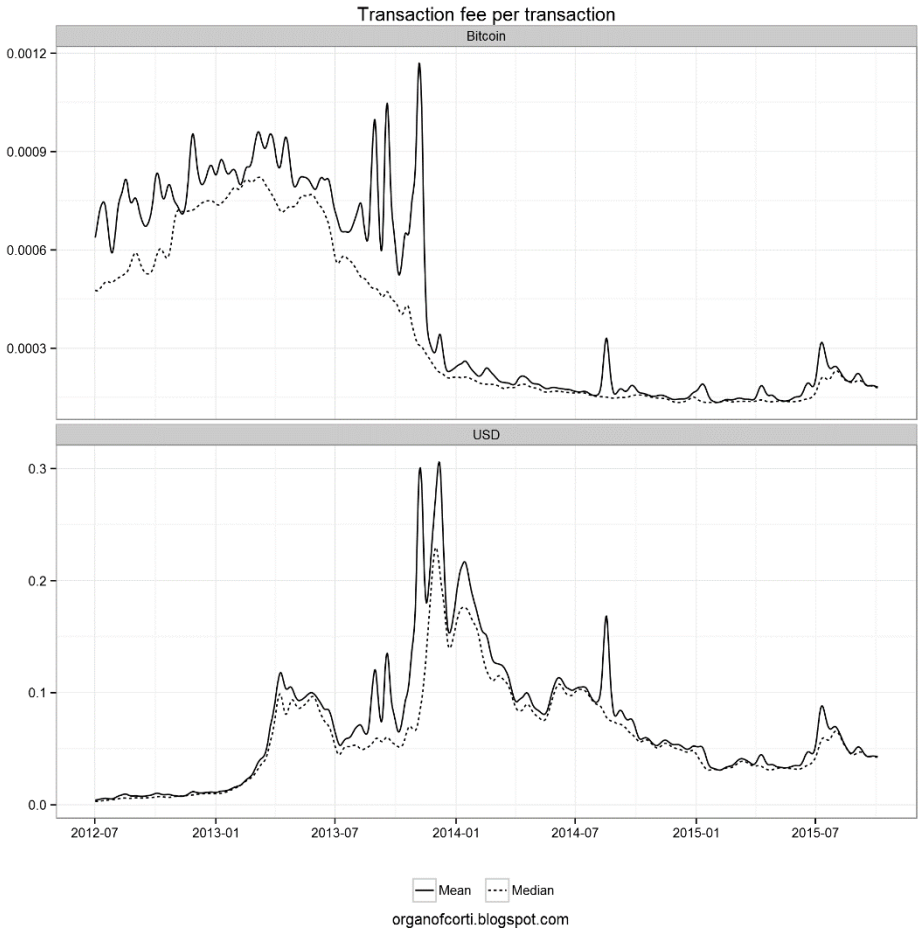

Source: Organ of Corti

The chart above illustrates transaction fees per transaction collected by miners denominated in both bitcoin (BTC) and USD since July 2012.[84] The recent uptick is largely related to multiple "stress tests" and "flood attacks" in which groups of individuals have sent significantly more than the average amount of transactions – which require a fee – to test the consequences of continuous "full blocks" on the network itself.[85] Much of this traffic is largely spam (e.g., "long chains") and does not represent actual commerce or trade.[86]

In Ethereum, this issue has been recognized, and one way the developers are attempting to mitigate the issue of disproportionalism is by charging fees for resource usage:[87]

> The intent of the fee system is to require an attacker to pay proportionately for every resource that they consume, including computation, bandwidth and storage; hence, any transaction that leads to the network consuming a greater amount of any of these resources must have a gas fee roughly proportional to the increment.

Due to its recent release, it is unclear at this time if mandatory resource usage fees on the Ethereum network will succeed in ameliorating this free-rider problem, or if it will follow the

same pattern Bitcoin has, and potentially lead to resource exhaustion attacks. To what extent there is a workable solution, based on (Luu et al. 2015), remains an open question.[88]

Source: Bitnodes

Secondly, as shown in the chart above, node levels have continued to decline.[89] Over the past year alone it has dropped about 20%.[90] And since March 2014, the drop has been more than 40%.[91][92]

Similarly, to the point addressed in (Rosenfeld 2012), if we are to define a "miner" as an entity that actually creates blocks (such as a mining pool) there are roughly the same amount of mining pools today (roughly 20) as there were in December 2012.[93]

While there is an open debate over how many nodes are needed or necessary to maintain pseudonymous decentralization, it bears mentioning again the key assumption that adding watermarked tokens to the network would lead to an increase in fees to nodes. As (Teo 2015) notes, this has yet to occur, though it is possible that user behavior will change for some unknown reason.[94]

(Rosenfeld 2012) briefly explores an assumption that links to a demerit of colored coins in relation to "blockchain bloat":

> The inclusion of extra burden of colored coin transactions can supposedly bloat the blockchain, increasing the cost of running a node. However, scalability is an issue for any blockchain even with only native transactions, and is an issue which should be solved rather than worked around. Every transaction carries the marginal cost of being received, verified and stored by every node on the network. These are all commodities, and transaction fees exist in part to pay for them. With a properly chosen system for the distribution of fees to nodes, any additional transaction makes it that much more

> lucrative to run a node, maintaining the balance of the number of nodes and their profitability. In fact, since new economic activity can bear above-cost fees, it can serve to actually strengthen the network. The overall level of fees is a protocol-level decision to balance the number of nodes with the reduction of friction; whatever level is chosen, additional layers can be used to accommodate transactions with a lower value than the threshold (for both the native currency and colored coins).

While bloat has arguably occurred via "long chains," the general assumptions above have not borne out. As explored in (Swanson 2015a), "long chains" is a catchall term used to describe the relative rapid velocity of transactions and is typically represented by on-chain casino payouts and tipping.[95]

Most importantly, there has been no mechanism created to sustainably pay nodes beyond altruism.

In fact, one of the fundamental problems of deciding who to pay, how often and how much is that the gossip network can be deceived via Sybil attack. By spinning up nodes on a cloud provider, an individual can create hundreds of fake nodes. One such incident occurred on August 25, 2015, in which 3,000 additional nodes were temporarily added to the Bitcoin network. This is an ongoing issue that the Bitcoin XT fork is also witnessing.[96]

Incidentally, proof-of-work was one of the solutions to make this type of spoofing uneconomical. Thus, as a counterfactual to (Rosenfeld 2012), it appears that one of the only ways to appropriately reward validators or "nodes" on a public blockchain is to first require the submission of proofs-of-work, which in practice brings the issue full circle. Aside from Bitfury and KnC, few participants are currently interested in investing in the capital needed to operate *both* a fully validating node and proof-of-work equipment.

Furthermore, (Rosenfeld 2012) had the foresight, however, to consider issues surrounding "protocol-level decision" which now envelops the community in the form of the "block size debate" (discussed in Appendix A). Two proposed solutions include duplex micropayment channels (Decker and Wattenhoffer 2015), as well as the Lightning Network, a "layer" marketed as having the ability to mitigate scalability limitations of Bitcoin and Bitcoin-like networks.[97]

Why are validators on the network unable to detect watermarked assets in order to levy higher fees?

Aside from mining pools that use Luke-Jr.'s software (such as Eligius), miners in general currently have no way to distinguish a watermarked transaction from any other transaction. Recall that at one point Eligius and a few other pools actively blocked transactions from their mempool that originated from casinos, such as SatoshiDice.[98]

Consequently, mining pools (and farms therein) have no incentive to destroy additional units of capital (energy) to protect against block reorganizations (reorg) and double-spending for these

watermarked tokens, mainly because they receive no additional revenue to do so.[99] This results in a top-heavy network.

For example, if Apple Inc. decided to issue all of its shares onto the Bitcoin blockchain via a metacoin such as Counterparty, they would have inadvertently created a top-heavy security vulnerability. Will Bitcoin security suffice to keep the Apple shares trading market secure? [100]

In this case, miners are securing not the metacoins themselves but securing against block reorgs and double-spending. Therefore, what is relevant is the cost of a reorg (i.e., the amount of capital an attacker has to consume) versus the value of the metacoin that is being unwound (and not the market cap of the metacoins itself).

If the cost of a reorg is less than the reward an attacker gains from doing the reorg, then the attacker could go ahead and reorg profitably. The relevant metric is the amount of money the attacker can make by unwinding trades, not necessarily the total value of the metacoins.

However, the money made by unwinding trades is directly correlated with the market value of the color and is just one possible payoff vector.[101] Why? Another point of doing a nefarious reorg is double-spending. In this case, an attacker would sell their Apple shares, convert the proceeds to cash and do a reorg pretending the sale never happened.[102]

Thus in the long run, miners are probably not destroying enough capital to ultimately secure against reorgs and double-spending of colored coin and metacoin assets, making the entire network less secure.[103,104]

Section 2.2: Turning it to eleven

This section looks at how: (1) Bitcoin mining pools are currently not color aware; (2) Making the Bitcoin network color aware creates security risks; (3) Even if Bitcoin mining pools were color aware, security cannot be increased due to its inherent assumption that all bitcoins are equal.

So we should ask: does adding exogenous value raise the risk of a double-spend? Perhaps it does because more hashrate is required to secure a proof-of-work blockchain that is attempting to guard the finality of off-chain assets. Yet, as per (Luu et al. 2015) there is no automatic mechanism to reward this additional labor leading to a possibility of having to remove some scripts altogether – or, in the case of Ethereum, a potential "verifier's dilemma" which then can lead to a fork.[105,106]

Yet if a pool becomes "color aware," so that it could be proportionally rewarded, it then opens up a new vulnerability: transactions could then be more easily singled out and censored. A model of one scenario, called the hold-up problem, can be found in (Teo 2015) and Appendix B.

Is there another way that pools can identify colored transactions? The OP_RETURN function of the Bitcoin network is one way that watermarked tokens could be identified by validators. However, use of OP_RETURN has been met with controversy.

For instance, in March 2014, when Bitcoin Core version 0.9 was released, there was a prolonged public dispute between Counterparty developers and several Bitcoin Core developers including Jeff Garzik and Luke-Jr. over the intended use of OP_RETURN. As noted in section 1, this is an op code that several colored coin and metacoin developers now commonly use as a type of data store.[107]

While the original debate circled around the proposed versus implemented store size (40 bytes versus 80 bytes), most of the discussion dovetailed in disagreements about what type of information was "allowed" to be stored in the first place.

As discussed in (Swanson 2014), while there are other ways to insert data permanently into the blockchain, following the finalization of 0.9 bitcoind, based on an explanatory post by Gavin Andresen, a Bitcoin Core developer, some of the community considered the enablement of OP_RETURN to be a feature, where the Bitcoin blockchain would be used as some kind of universal data store.[108,109] Yet as observed in the contemporary social media forums, other developers argued that OP_RETURN was not intended to be used as a general data store function and that it was to be used solely for encrypted keys (specifically ECDSA).

One solution to the impasse was proposed by Luke-Jr., who also serves as a consultant to the Eligius pool: enable pool software to filter and censor transactions that used OP_RETURN as a generic data store.[110] The dispute ended as soon as it began, but without any clear definitive "ruling."

A year later, in early 2015, OP_RETURN was increased from 40 bytes to 80 bytes, but the debate surrounding such filtering techniques still continues.[111] These governance challenges, over who ultimately is the "decider," are explored later in Appendix A.

What other limitations does "color blindness" have on the current pooling software and on the network?

This issue was raised by (Cohen 2015), who noted that since proof-of-work as used in Bitcoin and Bitcoin-like systems cannot offer *additional* levels of security, the proof-of-work knob cannot be turned to the proverbial "11":

> The other problem is it puts a limit on the amount of security the system can have. The waste gives the system the amount of security that went to it but it's limited to that; it won't have more than that. It would be really nice to be able to have more security than that. [112]

What if colored coins were obfuscated and miners remained "colored blind?" What impact would a heavily obfuscated colored coin system have on Bitcoin?

For instance, Bob could design a protocol that commits to metadata in the PAY_TO_SCRIPTHASH (P2SH) rather than in the OP_RETURN. Thus only clients "in the know" will see the color and not the miner. It can be done similarly to stealth addresses but other designs are possible.[113]

On the face of it, Bob would likely have to reveal the script when he collects a pay to script hash (P2SH) payment. But ignoring the cat-and-mouse game of obfuscation and heuristics used to uncover what is really taking place, several orders of magnitude in disproportional value being added to a public Bitcoin-like network would likely lead to governance issues as explored in Appendix A and potentially even a hold-up problem as explored in Appendix B.

In addition, if such a "decoloring" proposal was brought up, it could also face similar backlash as the OP_RETURN changes did previously and "block size debate" currently illustrates.

For example, the authors of (BitFury 2015a) speculate that one alternative to increasing block sizes is restricting certain types of transactions such as colored coins but also noted it could be contentious:

> Compared to a block size increase, transaction restrictions have an advantage: they constitute a soft fork (i.e., blocks mined by upgraded software with more restrictions would be accepted by non-upgraded software), whereas a block size increase implies a hard fork (i.e. change to the Bitcoin protocol that is not backward compatible). On the other hand, this solution to transaction throughput could cause a backlash among Bitcoin users and developers building applications on top of the blockchain.[114]

The debate surrounding block sizes is discussed later in Appendix A.

Section 2.3: What is merged mining and why it is considered important

Can merged mining stave off or provide additional security to these public ledgers? (Rosenfeld 2012) asserts that:

> Different cryptocurrencies work best if they pool their hashrate, as each of them enjoys the security of the combined hashrate. This can be achieved with merged mining; however, this still requires some amount of hashing native to the currency, as well as support of some miners from the host blockchain (such as Bitcoin). Smaller issues with few users will find it difficult to enjoy protection from hashrate-based attacks. Also, the host blockchain will not benefit from the native hashing of the guest blockchains.[115]

Merged mining is frequently described as a way of productively reusing the existing capital base of the mining network. Definitionally, **merged mining** (sometimes called auxiliary proof-of-work or AuxPOW) allows a mining pool (a block maker) to simultaneously mine for more than one blockchain.[116] That is to say, once a specific AuxPOW is added to mining pool software, it enables miners to produce valid hashes for other blockchains.

In theory it sounds like a viable option to extend the life or increase the security of a proof-of-work network but in practice, if it costs miners nothing to merge mine, then it also costs them nothing to attack the merged chain.

Relying on and trusting in goodwill, charity or altruism of a labor force (such as through hypothesized "assurance contracts") is the direct antithesis of the game theory baked into Bitcoin itself where it is assumed that there can and will be adversaries and incentives were designed so that way the network is not only endogenously self-sustaining but participants are motivated to build only on the longest chain.[117] As noted in (Goodman 2014), "A robust currency should not need to rely on charity to operate securely."[118]

The first known chain to undergo this AuxPOW process was Namecoin, which began merged mining with Bitcoin in October 2011.

Contrary to the claims in (Bitfury 2015b) we have empirically seen (below) how Namecoin's hashrate has diverged over the past year and how it now consistently represents less than one-third of Bitcoin's.[119]

This is due to at least 2 reasons:

1) not all Bitcoin pools support AuxPOW (merged mining) with Namecoin

2) also due to the Namecoin block reward halving that took place in mid-December 2014

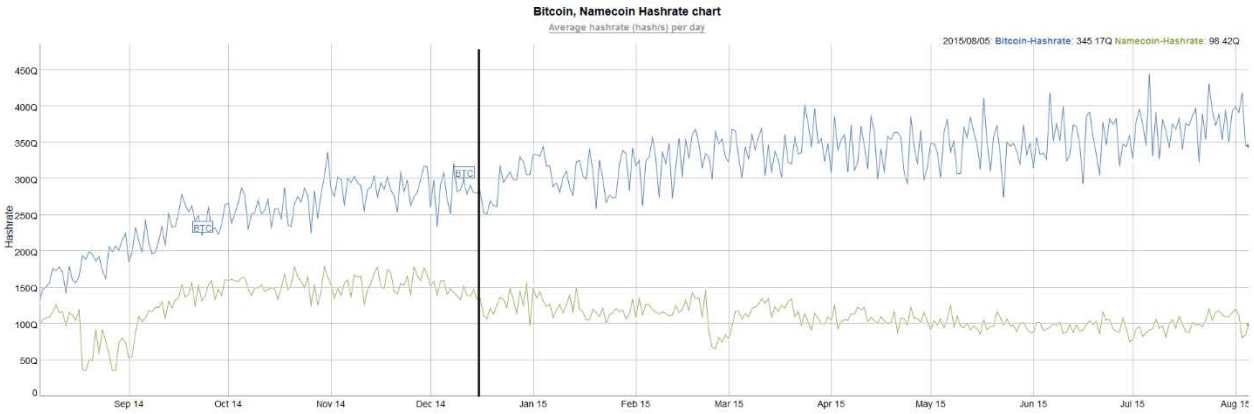

Source: Bitinfocharts.com

In the chart above, Namecoin's hashrate declined from around 130 petahashes in mid-December 2014 to roughly 100 petahashes as of this writing.[120] The black vertical line shows when the "halvening" (block reward halving) took place: when miners began receiving 25 namecoins per reward instead of 50.

Why did it decline? In contrast to the popular narrative, after the halvening there was in fact no doubling in namecoin market value because the market had already priced the future block halving into present day prices, so the reduction in block rewards acted like a pink slip to some marginal miners.[121]

Thus a challenge presents itself when this seigniorage subsidy is halved, a structural attribute of most cryptocurrencies. With Bitcoin, the subsidy is reduced by 50% every 4 years (or every 210,000 blocks).

While this issue typically remains hidden and muted when cryptocurrency values appreciate and rise, in the long run, continual halvings disincentivize laborers from providing security and utility to the network. There have been several cryptocurrencies whose labor force dissipated after their profitability period was over – most notably with Auroracoin and Feathercoin – and as a consequence the network became less secure.[122,123]

One such popular token that faced this dilemma last year was Dogecoin, which – as shown in the chart below – lost 20-30% of its security force every 2 months due to accelerated halving.

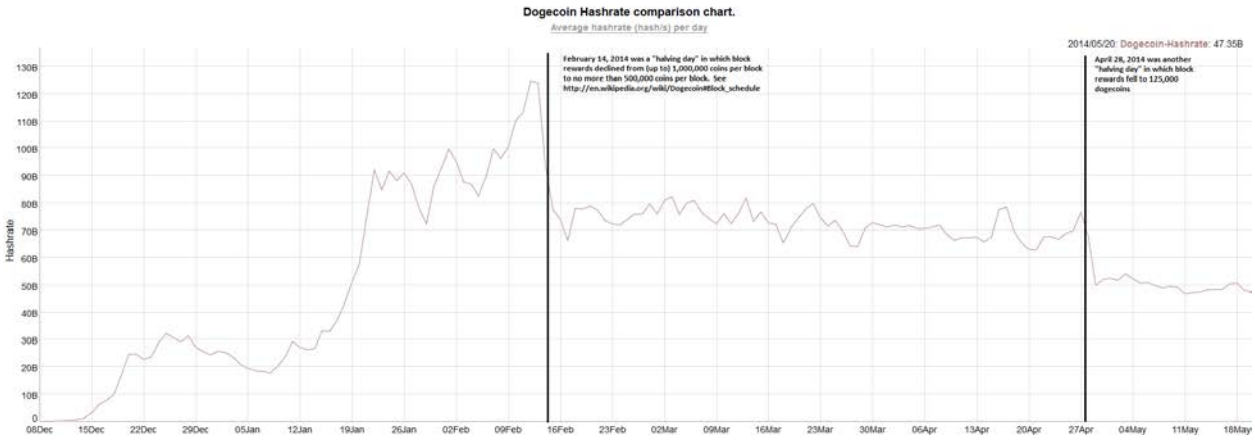

Source: Bitinfocharts.com

How to prevent or reverse this decline? While there were multiple potential solutions Dogecoin developers could adopt, incorporate or migrate to, on September 10, 2014, the lead developers of Dogecoin "switched on" the AuxPOW feature and convinced several Litecoin pools to merge mine as well.[124]

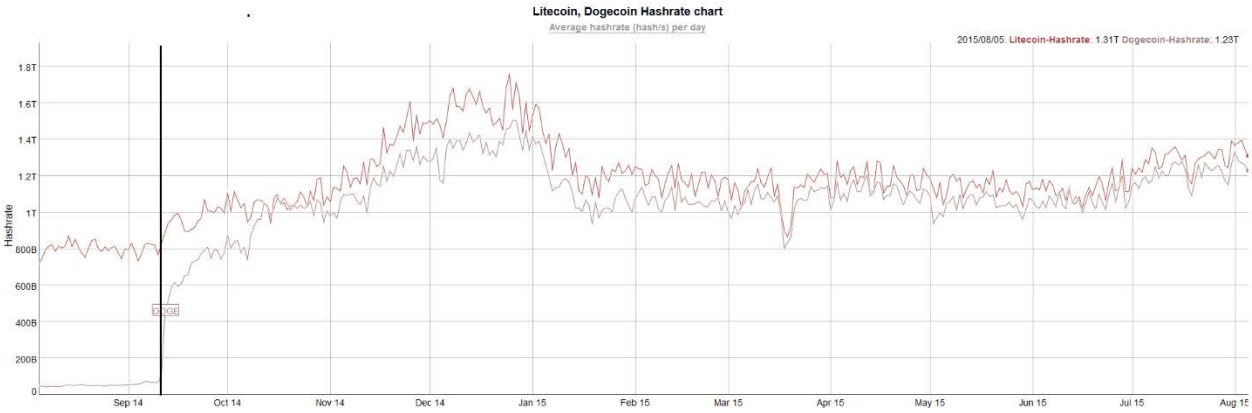

Source: Bitinfocharts.com

The black vertical line on the chart above shows the dramatic improvement Dogecoin had, jumping up to a hashrate roughly equivalent to Litecoin.[125]

Yet this phenomenon, as noted in (Saberhagen 2013, Courtois 2014), could be self-terminating in the long run as the Litecoin network itself may shed security due to additional halvenings, losing hashrate and therefore becoming less secure.[126,127] For instance, in August 2015, Litecoin underwent a halvening, resulting in a temporary 20% decline in hashrate over the course of several days before returning to its previous average rate.[128]

Why did the hashrate return to its moving average? One theory recently retold by Charlie Lee, creator of Litecoin, is that certain Litecoin hashing farms in China now allegedly have nearly zero electricity costs due to their owners' unique social connections with power plant operators.[129] At the time of this writing, the Litecoin hashrate is still about 30% lower than its all-time high, which means there is usable equipment that has been sidelined for unknown reasons.[130]

Is this analysis merely an academic exercise? No.

On September 12, 2015, an identity-focused startup called Onename announced that it would be migrating its service from the Namecoin blockchain to the Bitcoin blockchain.[131] Why? Because for several back-to-back months, one mining pool (F2Pool) continuously controlled more than 51% of the network hashrate and "could reorganize the blockchain, censor name registrations, deny updates, and squat names as they expire."[132]

It bears repeating that public networks such as Bitcoin do not simply take care of themselves. Maintaining miners and code development require the consumption of real resources.

Thus, we are left to wonder what the sufficient, sustainable incentives for proof-of-work cryptocurrencies to continue providing security are. There has been lots of idle speculation from large promoters and investors of Bitcoin companies, but very little research by Bitcoin companies beyond posturing on social media and at conferences.[133]

This is currently being modeled by a variety of academics. For instance, in (Teo 2014), the model results in a monopoly mining pool scenario due to future block reward halvings and a lack of increase in aggregate fees.[134]

Section 3: Traceability and permissionlessness

During research over the past six months related to distributed ledger projects which provide asset issuance and tracking, one of the common objections financial institutions have had with using a public blockchain revolves around the lack of knowledge of who is processing the transactions.

In the US, financial institutions are required by certain agencies, such as the Office of Foreign Assets Control (OFAC), to meet and fulfill certain requirements. OFAC is a US Department of Treasury agency that administers and enforces economic and trade sanctions based on US foreign policy and national security goals against targeted foreign countries and individuals. This includes checking the OFAC Specifically Designated Nationals (SDN) list for individuals and organizations with whom US citizens and residents are prohibited from doing business. Thus, in order to comply with OFAC requirements, financial institutions must conduct Know Your Customer (KYC) on the parties with whom they transact.

While OFAC has not ruled on this point, this obligation at least arguably would apply to transaction fees paid to miners.[135] If a financial institution sends a transaction fee to a Bitcoin miner that is in a sanctioned country or is known to facilitate illicit transactions, then the institution could be held liable for breaching those regulations. Because mining pools and farms involved in the Bitcoin hashing process are not gated or vetted, it cannot be assured that the miner is not in a sanctioned country (e.g., Russia, Syria, North Korea).

But aren't networks like Bitcoin decentralized to the point where it is difficult to identify a specific miner or block maker?

In theory, yes. Over time, and due to incentives to centralize hashrate into pools, block making has converged into a quasidynamic group of just under 20 "miners."[136]

Recall that while there may be around 5,000 nodes on the "gossip" network, propagating blocks around the world, the process of mining – of transaction selection and verification – is what actually matters in terms of security via proofs-of-work.

While all nodes may have the entire chain history, in practice it is only *one* pool that creates blocks. And over time, beginning in the winter of 2010 and spring of 2011, these pools began self-identifying themselves to the public.[137]

Consequently, we do have a general idea of where many pools and farms are located.[138] Common jurisdictions include:

- several provinces in China (from where ~55% of the hashrate currently originates),
- Iceland (where both Genesis Mining and BitFury have operations),[139]
- Finland and Republic of Georgia (out of where BitFury primarily operates),
- Ukraine (GHash.io),
- Sweden, Iceland and Finland (KnCMiner),
- Czech Republic (Slush) and
- Washington State and Arizona (where Mega Big Power and purportedly 21inc reside).

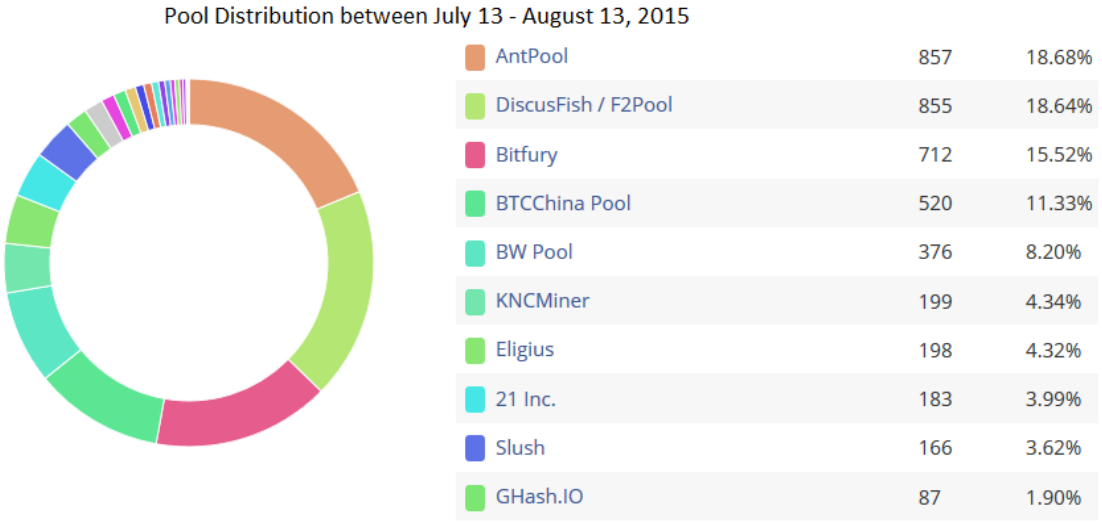

Distribution of block making over 30 days. Source: Blocktrail

In addition, young startups such as Blockseer, Chainalysis, ScoreChain, Coinalytics, Elliptic and Sabr are building AML services that enable organizations and institutions to model cryptocurrency funds back to their origin.[140] Which leads to an interesting phenomenon.

Based on the current trend, within a few years, the edges of the Bitcoin network may become fully KYC'ed and the validators simultaneously could become KYM'ed – either through outright top-down mandates, or forks (such as Bitcoin-NG, Bitcoin XT, XT 2, XT 3 and so forth) or through data analytics. These in turn reduce the primary advantages of a permissionless network: permissionlessness.[141]

Arguably the most accurate term to describe this phenomenon is **permissioned-on-permissionless (P-o-P)**.[142]

Recall that a permissioned system is one in which identity for users is whitelisted (or blacklisted) through some type of KYB or KYC procedure. It is the common method of managing identity in traditional finance.[143][144]

In contrast, a permissionless system is one in which the identity of validators is pseudonymous or even anonymous. Bitcoin was originally designed with permissionless parameters and goals,

although as of this writing, as described above, many of the on-ramps and off-ramps for Bitcoin are increasingly permission-based and, in some cases, act similar to existing banks and payment service providers.

A P-o-P system is a hybrid in which the on-ramps and off-ramps of the networks are gated via KYC and AML policies – as are some of the validators themselves – yet some of the block making and transaction selection is still done pseudonymously. Thereby, the high marginal costs of proof-of-work remain, as visualized below.

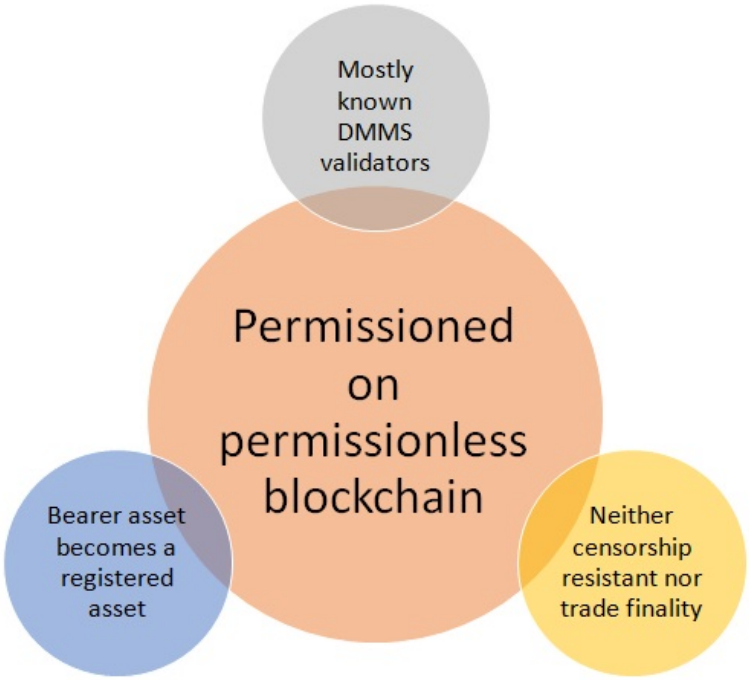

Figure 1: Permissioned on Permissionless (PoP)

This is to say, while the gating process to become a validator is still relatively permissionless (in the sense that no single entity authorizes whether or not someone can or cannot create proofs-of-work), the fact that pools are self-identifying is a bit ironic considering the motivations for building this network in the first place: creating an ecosystem in which pseudonymous and anonymous interactions can take place. If you know and trust the counterparties, you could use other existing, less capital intensive methods for securing transactions.

The second bucket, neither censorship resistant nor trade finality, refers to the fact that large venture-funded companies like Coinbase or Circle not only require identification of its user base but also censor their customers for participating in trading activity that runs afoul of their terms of service.[145]

Technically speaking, hurdles to *on-chain* "trade finality" refers to bitcoin transactions not being final. Due to a block reorg, a longer chain can always be found undoing what you thought was a confirmed transaction. This has happened several times, including notably in March 2013.[146,147]

Of the $900 million that has been invested in Bitcoin and Bitcoin-related startups over the past 3 years, more than 60% (or $540 million) of the funding has gone to startups that are building such P-o-P systems.[148] This figure draws directly from more than two dozen Bitcoin-related startups that have received a Series A or higher, nearly all of which now require mandatory KYC of its users. Thus the total amount of P-o-P is likely higher as this figure above does not take into account seed stage companies which also require KYC, KYB and KYM.

But isn't Bitcoin mining still permissionless?

Perhaps this was the case in 2009 and 2010. However, many mining manufacturers and pools such as MegaBigPower and 21inc, are now collecting KYC/KYM on their users.[149]

For instance, manufacturers and pools such as KnC Miner and BitFury are unopen to the public (they do not operate a publicly accessible pool), yet because they raised outside funds from venture capitalists, the operators have effectively submitted KYM information to a variety of authorities. Notable exceptions to collecting KYM information include most Chinese-based pools, except BTCC Pool which does require KYM.

What did it look like in 2009?

Recall that one way to describe the process of mining (or block making) is a dynamic-membership multi-party signature (DMMS), which was defined in (Back et al. 2014).[150,151]

What does DMMS mean? Originally there was no gating or authorizing process to enroll for creating and submitting proofs-of-work: theoretically, validating Bitcoin transactions is permissionless.

"Dynamic-membership" means there is no fixed list of signatories that can sign (i.e. anyone in theory can). "Multi-party" effectively means "many entities can take part," similar to secure multiparty computation.[152]

In other permission-based terms: producing the correct proof-of-work, which meets the target guidelines, permits the miner (block maker) to have full authority to decide which transactions get confirmed.

In other words, other than producing the proof-of-work, miners do not need additional buy-in or vetting from other parties to confirm transactions onto the Bitcoin blockchain. The "signature" on a block is ultimately signed by *one* entity and does not, by itself, prove anything about how many people or organizations contributed to it.

The original permissionless system is visualized thusly:

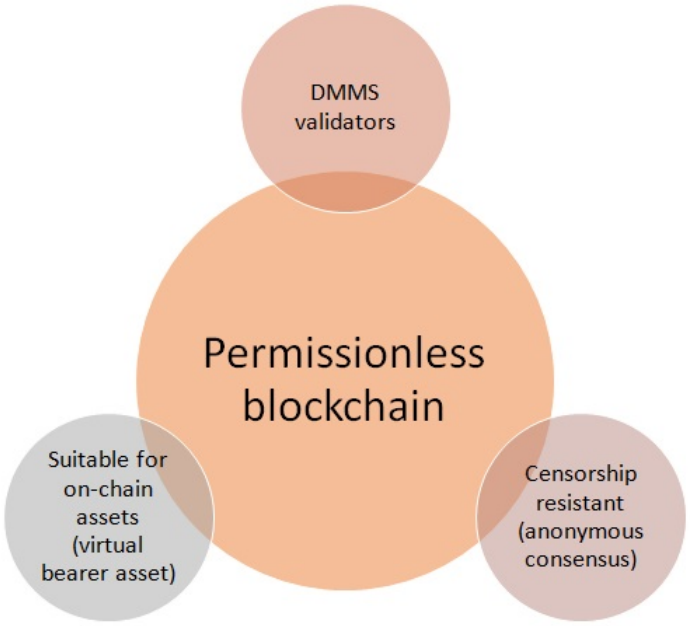

Figure 2: The idealized Bitcoin network circa 2009

However, as noted by the empirical trend towards P-o-P, this strawman model in Figure 3 is not how the system or ecosystem necessarily behaves *in practice*.

What impact does permissioned-on-permissionless have on startups?

In early April 2015, MegaBigPower and Bitmain reported that Coinbase, a large venture capital-backed Bitcoin intermediary, "has asked for detailed information regarding their mining facilities, requesting time-stamped photographs and videos of mining facilities and details about the origin of the hardware the companies operate."[153]

At the time of this writing, the pool operated by MegaBigPower represents around 1% of the network. Bitmain is a Beijing-based company that manufacturers hashing equipment used by many farms around the globe, accounting for around 55-60% of all hashrates.[154]

Section 3.1: Government sanctions

Because colored coins and metacoins depend on a nominal amount of bitcoins to operate, what happens in the event that the underlying bitcoins, or the miner fees associated with the block that contains the watermarked asset, are sent to a sanctioned individual or organization?

The new documentation requirements for MegaBigPower and Bitmain noted in the previous section may be related to a legal issue that Coinbase encountered two months earlier in February 2015.

According to *Washington Free Beacon*, an independent news site, during Coinbase's Series C fundraising effort, the startup distributed a pitch deck that included a known use-case that proof-of-work networks have: censorship-resistance.[155]

The deck, made available by *Washington Free Beacon*, noted this by stating one advantage for using Bitcoin is that it is "Immune to country-specific sanction (e.g., Russia-Visa)."[156]

What was the Russia-Visa sanction story?

According to *Reuters*, on December 26, 2014:

> The world's two largest credit and debit card companies, Visa Inc and MasterCard Inc , said on Friday they could no longer support bank cards being used in Crimea, following U.S. sanctions imposed earlier this month.[157]

It is well-known that one of the use-cases of public blockchains such as Bitcoin is that it could re-route bitcoins around censored regions of the world.

In fact, according to the abstract of the original white paper, one of the primary motivations for creating Bitcoin in the first place was to enable; a "peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution."[158] In other words, censorship-resistant e-cash.

This story may not seem noteworthy since several custody wallets and exchanges serving US customers do not fully operate in compliance with US federal regulations (e.g., filing SARs), but as *Washington Free Beacon* noted later, in the same pitch deck Coinbase claimed to have strong relationships with regulators whom reportedly were unhappy with the specific highlighted use-case.[159][160]

This illustrates the struggle and ongoing debate between permissioned and permissionless endeavors which simultaneously intersect with both pseudonymous and KYC environments. Many startups that raise outside funding from venture capital empirically have attempted to become more compliant with existing laws and regulations. But in doing so, by establishing mandatory KYC and AML documentation, it runs counter to the original utility propositions of these pseudonymous systems.

For instance, financial institutions serving US-based customers are required to check source-of-funds with the OFAC list. And according to (Hoegner 2015):

> U.S. persons are prohibited from engaging in prohibited transactions, or trade or financial transactions defined by the sanctions program, unless authorized by OFAC or expressly exempted by statute.[161] Penalties for engaging in a prohibited transaction are severe and typically based upon "strict liability," i.e., it does not matter whether a person knows that she is conducting or facilitating a prohibited transaction with a prohibited person. The fact of the conduct or facilitation of a prohibited transaction is sufficient. It is conceivable that Bitcoin could be used to transfer funds to sanctioned countries, groups, or individuals, and thus evade sanctions.[162]

While it has not been tested in court, one possible interpretation is that financial institutions should fulfill a "source of funds" request with miners, which results in a Catch-22: if you *cannot* identify miners, then you probably should not send them transactions, because they could be using bitcoin fees for illicit purposes, or may reside in a sanctioned country.[163,164]

If you can identify the miners, you might as well use a permissioned ledger or a database, because you end up with an increasingly narrower set of KYM'ed pools, thereby defeating the *raison d'etre* of paying for a pseudonymous, censorship-resistant network.

Why is this worth re-emphasizing?

For example, in two recent papers (Bitfury 2015b, 2015c), Bitfury – a large Bitcoin mining company based in the Republic of Georgia – endorsed the formalization of "trusted" miners into a set of fixed transaction processors.[165]

However, having a "trusted" entity run a proof-of-work mining farm is self-defeating with respect to maintaining pseudonymity on an untrusted network (which were the assumptions of Bitcoin circa 2008). It is a contradiction.

If all miners are accredited or KYM'ed then you do not have a Sybil issue and instead are now operating a very expensive trusted network. This directly conflicts with the dynamic part of a DMMS and with the security assumptions (and costs) of the original experiment itself.

Consequently if the Bitfury narrative gained traction then in practice financial institutions would still have to trust a fixed set of mining farms which is merely the recreation of payment service providers, yet without the supervision of FinCEN or other regulatory bodies.[166][167]

But isn't this making a mountain out of a molehill?

At this time, miners do not ever have control of the private keys.[168] However, they do have the option not to execute, block and potentially reverse transactions (through a block reorg) which impacts any virtual currency, colored coin or metacoin using the underlying network.[169]

How does this take place?

The mining process preserves the separate identity of the property (the unspent transaction outputs). And therefore the chain of title remains complete, as does the paper trail. Recall that there is a distinction between as acting as a facilitator and acting as intermediary. An intermediary takes title, whereas an agent, or one type of facilitator, does not.[170] They may take possession, but title remains with the principal (the Sender).

In the case of Bitcoin, it is assumed that title goes from Sender to Receiver and that miners sit in the middle, but does not take title (i.e., serves as an agent of the Sender).

While the miner does not hold the title itself, that is how agency works: they act in the name of principal.[171] One such analog is armored cash transport cars, though there is also a regulatory or statutory exception regarding communication and network access. [172]

Whether or not this is a legal concern for institutions serving US customers is unclear at this time. The real point is that there is OFAC risk in dealing with a public blockchain where transactions are secured by unknown miners.

The European Banking Authority (EBA) in their July 2014 report also highlighted these concerns surrounding sanctions:[173]

> VC [Virtual currency] transactions are not recorded and are anonymous, global and irrevocable. Also, decentralised VC transactions are not dependent on entities on which financial sanctions and embargoes could be imposed. As a result, it is difficult for governments and international governmental organisations to enforce financial sanctions or embargos against other jurisdictions, for example to further humanitarian objectives. The priority of the risk is high.

In point of fact, public blockchains such as Bitcoin are not economically irrevocable. Mutability and finality of transactions are *only* probabilistic (Garaye et al. 2015): the probability of a history-reversing attack rests on economics, not technology.[174] This creates a fundamental problem with financial systems that need definitive title transfer and settlement finality.[175][176]

As explained in (Sams 2015d):

> The financial system and its regulators go to great lengths to ensure that something called “settlement finality” takes place. There is a point in time in which a trade brings about the transfer of ownership – definitively. At some point, settlement instructions are irrevocable and transactions are irreversible. This is a core design principle of the financial system because ambiguity about settlement finality is a systemic risk. Imagine if the line items of financial institution’s balance sheet were only probabilistic. You own X of Y with 97.5% probability. That is, effectively, what a proof-of-work based distributed ledger gives you. Except that you don’t know what the probabilities are because the attack vectors are based not on provable results from computers science but economic models. Should a settlement system be built on that edifice?[177]

In other words, a court will only recognize a ledger that it deems authoritative. It may not be tolerable to have probabilistic finality for firms dealing with high-value off-chain titles.[178]

Shouldn't banks just purchase mining equipment to prevent this from happening?

For example, Fred Wilson, a venture capitalist who has invested in multiple Bitcoin companies, is a proponent of banks buying mining hardware. He recently downplayed the legal and pseudonymity concerns at financial institutions:

> One concern I hear, though, is that banks like to know who is managing their infrastructure and they are uncomfortable with miners they don't know, located in parts of the world that make them nervous, providing the transaction processing infrastructure for these applications being built on the blockchain.
>
> To me, that is the perfect reason for banks and brokerage firms to take a bit of their data processing infrastructure and point it to the blockchain and start mining it. They could even create a mining pool among the large money center banks. And it is relatively simple for a blockchain application to route its transactions to certain miners to process.[179]

While a shared ledger may be useful for banks in general, building a Bitcoin mining pool at a bank would simply add more costs and legal compliance issues without providing a similar amount of utility.[180] For instance, (Bitfury 2015b) proposed that banks spend $10 million for proof-of-work security: a service which is entirely unnecessary in a banking environment.[181]

In conversations with financial institutions, banks are uninterested in a censorship-resistant platform because they already know all the participants in their network, and are required by law to identify the counterparties of all transactions. Thus, proof-of-work is not only redundant, but would become a new cost driver.[182]

And if financial institutions had to bear the capital intensive burden of operating a mining pool to protect colored coin issued assets, as explored in (Teo 2015), they could avoid the costs altogether by simply using permissioned ledgers or databases to track assets instead.

To resolve this concern over pseudonymous validation, the EBA proposed several potential regulatory approaches in the long-term.

One such recommendation called for the creation of "an entity that is accountable to the regulator would need to be a mandatory requirement for a VC [virtual currency] scheme to be regulated as a financial service and for it to be allowed to interact with existing regulated financial services."[183]

And how would this be done? Through a term they call a "scheme governance authority" or SGA:

> The entity would be called the 'scheme governance authority', which is a non - governmental entity that establishes and governs the rules for the use of a particular [virtual currency] scheme. It is a legal person, and is responsible for maintaining the integrity of the central transaction ledger, the protocol, and any other core functional component of the scheme.[184] The scheme governance authority would be required to comply with regulatory and supervisory requirements of various kinds to mitigate identified risks. [185]

One potential candidate is a system called TrustGuard, proposed by (Srivatsa et al. 2005.)[186,187] However, a more likely candidate for an SGA-compatible entity is most closely related to a "mintette" – a known, legally accountable transaction validator – which was a term created in (Laurie 2011).[188]

(Laurie 2011) explains the role of this contractually-obligated validator, coining the new name:

> Leaving aside temporarily the question of how servers in the central authority (I shall call these mintettes) come to agree a transaction log, once this is in place we are now in quite a nice situation. At any point, a client can request from a mintette the current location of a coin and get a coin record, a snapshot hash and a proof that the coin record was included in the snapshot hash. If it was paranoid, it could also request the entire state for that hash and verify that the coin was not included twice. With this in hand, it could go to other mintettes and verify that the snapshot hash was current, or at least recent. Any mintette that attempts to fib about the current agreed state will quickly be revealed by very cheap checks at other mintettes.
>
> If mintettes also sign their responses, then any such naughtiness can be published as proof that the mintette should not be trusted. In a similar manner to the coin state, mintettes could keep a current list of blacklisted mintettes, including proofs of their malfeasance.
>
> Furthermore, mintettes can verifiably enforce constraints on the system. For example, if we wanted to produce a new coin every ten minutes, as the Bitcoin system does, mintettes would reject transactions that violated this requirement.

This idea was later expanded by (Danezis and Meiklejohn 2015), creating a system called RSCoin which is similar to Fedcoin (Koning 2014; Andolfatto 2015; Brown 2015a; Sams 2015a; Winkler 2015), a hypothetical idea revolving around a truly digital currency (as opposed to a *virtual* currency) that is administered by a central bank on a distributed ledger and is legal tender.[189]

The motivations by (Danezis and Meiklejohn 2015) were:

> Current cryptocurrencies, starting with Bitcoin, build a decentralized blockchain-based transaction ledger, maintained through proofs-of-work that also generate a monetary supply. Such decentralization has benefits, such as independence from national political

control, but also significant limitations in terms of scalability and computational cost. We introduce RSCoin, a cryptocurrency framework in which central banks maintain complete control over the monetary supply, but rely on a distributed set of authorities, or *mintettes*, to prevent double-spending. While monetary policy is centralized, RSCoin still provides strong transparency and auditability guarantees. We demonstrate, both theoretically and experimentally, the benefits of a modest degree of centralization, such as the elimination of wasteful hashing and a scalable system for avoiding double-spending attacks.[190]

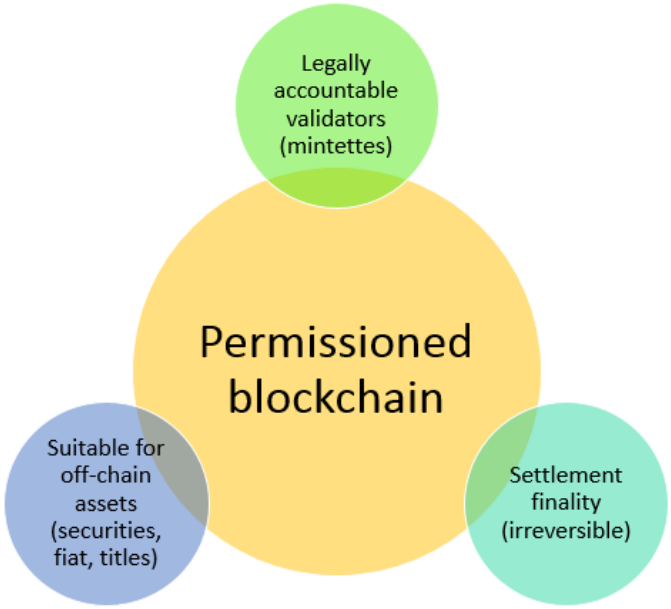

Figure 5: Permissioned blockchain

How is this visualized?

Above is a diagram first described by (Sams 2015b) and later discussed in further detail in (Swanson 2015b).[191]

So how does an enterprise or institution factor the concentration of mining pools into its risk analysis? What about their potential for pools to be manipulated externally?

To reduce risks, this question has prompted consideration of gated entry requirements to become a validating node on public blockchains, but will likely be poorly received by the decentralization community.

Section 3.2: Encumbered titles

In the seminal paper (Rogers 1990), the author looked at the history of negotiable investment certificates and why negotiability diminished over time.[192] Negotiability is a legal term referring to a document (such as a check or bill of lading or certificate). This enables the passing of ownership from one party to another, with the assurance that the new owner is protected from defects or liens from the previous owner.

With the advent of electronic record-keeping (electronic book-entry) and rise of computerized trading, the trend towards automation, digitization and dematerialization of financial instruments and titles has been ongoing for multiple decades.[193] This not only changed the landscape of who is involved in exchanging wares, but also what wares were involved.

For instance, the historical move away from bearer instruments was done so for a very practical reason.

According to (Rogers 1990):

> The demise of negotiability in the world of investment securities nicely illustrates the various causes that are gradually but surely leading to the extinction of the doctrine in all of its applications. First, negotiability doctrine rests on the assumption that the best way to transfer abstract rights is to embody them in pieces of paper and then physically deliver the papers from person to person. As the volume and velocity of trading increases, the requirement of physical delivery becomes an intolerable burden; once we pass from paper to electronic recording of financial relationships, delivery becomes a metaphysical absurdity.
>
> Second, it is by no means obvious that protecting purchasers facilitates the operation of the market. For every purchaser who is protected by negotiability doctrine there is an owner who is harmed. While I have neither the training nor inclination to endeavor a quasimathematical demonstration, a simple seat-of-the-pants observation will suffice. I never carry very much cash in the seat of my pants precisely because I know that if I lose it, it's gone. Indeed, one of the reasons that the Treasury ceased issuing securities in certificated bearer form seems to have been a concern that their negotiability which necessarily means vulnerability to theft-diminished rather than enhanced their attractiveness as investment vehicles.[194]
>
> Third, there is the phenomenon that may, in the end, finally seal the fate of negotiability for any form of investment or financial instrument. Any instrument that can be transferred by mere delivery is an instrument that can be transferred without any paper trail. The taxman and the policeman do not like that, and eventually they will prevail. It is no accident that cocaine deals are the only large dollar transactions in commerce that are settled in cash, nor that U.S. currency is no longer printed in large denominations.[195]

The third and final reason garners a lot of attention at Bitcoin-related conferences. A large percentage of Bitcoin adopters are involved because of various political motivations, which has manifested itself in the "Blockchain 2.0" world. While some distributed ledgers attempt to address all of these shortcomings, not all of them can do so in practice.

Watermarked tokens and decentralized cryptocurrencies, at first glance, appear to be new solutions to issue and transfer registered assets. For instance, (Rosenfeld 2012) explains that:

> With colored coins embedded in a host blockchain, any legal issue with the guest assets, such as uncertainty about security trade regulations or specific unsavory assets, might project into the host blockchain. This issue deserves more exploration, but it is the author's belief that the inclusion of multiple heterogeneous entities within the same blockchain will only reinforce its position as defying traditional legal approaches.[196]

Perhaps this narrative will eventually take place, but so far it has not. In fact, the opposite has occurred: nearly all of the funding that has taken place within the larger "blockchain" ecosystem requires some semblance of compliance with local and national regulations.

As has been pointed out multiple times, permissionless networks such as Bitcoin were designed to facilitate the movement of censorship-resistant cash, which was its priority. [197] One of the initial motivations for creating a public blockchain such as Bitcoin was due to a desire not to interact with financial institutions and governments.[198]

In contrast, the use-cases for permissioned ledgers does not include censorship-resistant cash, because it is difficult, although not technically impossible, to build censorship-resistant electronic cash on top of a permissioned network.[199]

What does this have to do with encumbered titles?

According to a recent compilation by Consumers' Research, approximately 1,739,530 bitcoins have been lost, stolen, hacked or scammed from users between August 2010 and March 2015.[200,201] This number is likely a lower-bound, as many users simply do not report their losses in a public forum. This represents roughly 12% of the bitcoin monetary base that is otherwise potentially encumbered.

In many cases, such as the successful December 2014 phishing attack on BitPay which resulted in the loss of 5,000 bitcoins, the title to this property is actually encumbered.[202] This has led to speculation that since many of these bitcoins are eventually intermixed and pooled with others, a large percentage of the collective monetary base does not have clean title. The implications of this can be far reaching for an asset that is not exempted from *nemo dat quod non habet*, the maxim of commercial law that one cannot transfer what one does not have.[203] Currency or legal tender, represents the sole unconditional exception to the *nemo dat* rule. Bitcoin however is not legal tender.[204]

How does this work? Since many users in general don't trust themselves with securing their own bitcoins, users have given (deposited) their private keys with a new batch of intermediaries that euphemistically market themselves as "custodial wallets" and "vaults." Many of these entities, in turn, comingle the bitcoins.[205]

Consequently, a non-insignificant portion of Bitcoin's money supply does not have unambiguous title. This potentially leads to problems in settlement finality, as legal title may be ambiguous.[206][207]

How does this impact bitcoins which underlie watermarked tokens? If the bitcoins that are being used to move a "colored" asset are encumbered, how does that impact the actual asset? Unfortunately, this legal issue has not been analyzed by courts at this time.

Can you have settlement finality on a public blockchain? Under existing laws, maybe not.[208,209]

Section 3.3: Uniform Commercial Code

If Bitcoin itself and "Blockchain 2.0" implementations such as watermarked tokens are unable, under the current Uniform Commerce Code (UCC), to provide settlement finality, how else does UCC impact virtual currencies?[210]

According to (Fogg 2015), while the Bitcoin network creates transparency in terms of where transactions have come and gone, it does not provide information on liens.[211]

Under the current UCC, bitcoins fall under the category of Article 9, as general intangibles.[212] As long as bitcoins are treated as such, no one can know for certain that a lien holder will not appear and claim ownership of encumbered tokens.

As noted by (Fogg 2015):

> A security interest in general intangibles like bitcoin continues, notwithstanding the sale, license or other disposition of the collateral, unless the secured party consents to the transfer free of its security interest, the obligations secured by the security interest have been satisfied or the security interest has otherwise terminated.[213]
>
> Thus, each time a bitcoin passes through the hands of an owner whose property is subject to a security interest in general intangibles that bitcoin becomes burdened with that security interest.
>
> Security interests in general intangibles are very common, especially for a business with a bank secured line of credit. It is the combination of bitcoins being general intangibles, security interests in general intangibles not being automatically released upon transfer

> and the ubiquity of security interests in general intangibles that creates the fatal flaw – bitcoins that are encumbered with security interests granted by one or more prior owners.[214]

And as described in section 3.2, because of how bitcoins have been comingled and absconded over the years, there "are probably more people with legitimate claims over bitcoins than there are bitcoins."[215]

Recall that in section 1, (Schroeder 2015) notes that the identification of ownership as it relates to the transfer of common stock is important. The author came to a similar conclusion as (Fogg 2015), noting that:

> The bad news is that bitcoin does not, and cannot be made to fit into, the U.C.C.'s definition of "money". If held directly by the owner, bitcoin constitutes a "general intangible". Unfortunately, general intangibles are non-negotiable. That is, unlike virtually every other category of personal property recognized by Article 9, once a general intangible becomes encumbered by a security interest, it can never become unencumbered even by transfer to a bonafide purchaser for value. This could greatly impinge on bitcoin's liquidity and, therefore, its utility as a payment system.

Remember that based on how the Bitcoin network was designed, knowledge of a private key (or credentials) is the sole method of controlling unspent transaction outputs (bitcoins).

This virtual currency is effectively a bearer asset that lacks any "information regarding existing security interests. This circumstance leaves acquirers of bitcoins with the risk that their bitcoins may be subject to security interests that diminish or eliminate the economic value of the bitcoins to owners."[216]

How can Bitcoin (and, conceivably, the platforms that rely on bitcoins, such as watermarked tokens) remove these uncertainties and deficiencies?

According to both (Fogg 2015) and (Schroeder 2015), one solution is for a securities intermediary (such as a custodian) to take custody of these financial assets as a type of "investment property" under UCC article 8.[217]

Whether or not a securities intermediary will attempt to do so is an open question.[218,219]

Couldn't the UCC simply be amended?[220]

Miles Cowan, a corporate attorney and co-founder of Ldger, a marketplace loan fintech startup based in New York City, has analyzed the impact UCC has on virtual currencies. In his view:

> Amending the UCC is not a trivial process, involving first getting a proposed amendment drafted by the Uniform Law Commission then getting the proposed amendment

> adopted by all 50 states (and the District of Columbia, the Commonwealth of Puerto Rico, Guam and the U.S. Virgin Islands). Given enough time we may get favorable bankruptcy decisions. Or a foreign government may authorize or adopt one or more cryptocurrencies as a medium of exchange.[221]

Why is this an issue for watermarked tokens such as "cryptosecurities?" It is a challenge because of disputes over conflicting ownership claims.

As (Schroeder 2015) hypothesizes:

> Presumably, most thieves who steal securities do so with the intent to sell them quickly before the theft is discovered. With a conventional certificated security, this might take several days so that the registered owner might have the practical ability to give an effective stop transfer order. However, one of the greatest advantages – and in this case the disadvantage – of a blockchain is its speed. By the time the registered owner discovered that a cryptosecurity was missing, the thief's sale would probably already have been settled so that it would be too late to give notice.[222]

What about ill-gotten performance of the watermarked financial instruments?[223] That is to say, in the event that sidechains, colored coins or other secondary protocols attached to Bitcoin begin issuing financial derivatives or even something as simple as a dividend, what happens in the event of theft?

For instance, if Bob's colored coin representing a stock is stolen by Alice, does the company issuing the dividend still pay it? If yes, then this could open up the company to lawsuits for knowingly paying a thief. If no, then Bitcoin (the network) ceases to function as a decentralized pseudonymous ledger, as the network would need to verify the identity, taking away the pseudonymity behind it.

Currently, there is no way to do this on Bitcoin itself today without a major code rewrite and change in the social contract which, as shown in the current block size debate, is a difficult coordination problem to resolve.

The legal issues surrounding virtual currencies such as Bitcoin, with respect to the Uniform Commercial Code, will likely be an area of continued research.

Section 3.4: FinCEN impact on watermarked tokens

Financial Crimes Enforcement Network, or "FinCEN," is an agency of the US Treasury Department that is responsible for enforcing the Bank Secrecy Act (BSA) and the anti-money laundering rules FinCEN adopted thereunder.

On August 14, 2015, FinCEN released a ruling very relevant to watermarked token platforms[224] FinCEN ruled that a company that offers brokerage services for precious metals and issues

Bitcoin-based digital "proofs of custody" for the precious metals it stores on behalf of its customers, is required to register with FinCEN as a money transmitters.[225]

Stretching back into 2013, there have been more than a dozen startups that have used colored coin projects and platforms such as Counterparty to effectively track embedded data that has been hashed or anchored into the Bitcoin blockchain, and which allegedly represents a fraction (or sometimes a whole piece) of ownership in some type of precious metal. In many cases, this is a gold or silver bullion.

While the ruling is directly applicable only to the subject company, it raises the question of whether other companies using Bitcoin or other virtual currencies to track and exchange other assets might also be required to register as money transmitters with FinCEN, and comply with the BSA's AML requirements, effectively turning them into a permissioned-on-permissionless platform.

According to Juan Llanos, an AML and risk management expert, the new FinCEN guidance may impact other companies as well, beyond those issuing proof-of-custody services:

> The broader statement is that issuing a certificate of ownership (whether paper, digital or even a statement) and allowing the unrestricted transfer of value is what makes a broker or dealer (or any other company, for that matter) a money transmitter.[226]

Over the past year a popularly cited use-case has been a tool called 'proof-of-existence' – a type of method that takes a hash of a document (such as a certificate) and embeds it onto the Bitcoin blockchain.[227] Whether a commodity-backed virtual currency is the equivalent of a freely transferable digital certificate in the eyes of FinCEN remains an open question.[228] Therefore it is also unclear if the scope of this guidance applies to any platform that enables digitally hashed document (such as patents) whom would then face the prospect of being a money transmitter.[229]

While it appears clear that law enforcement and regulators will continue enforcing existing money transmission laws and that individuals and groups who want to do illegal things with blockchains can't, it is relevant to highlight that when pseudonymity is removed from public blockchains then so too does some of the core functionality of a public blockchain.

For instance, in the event that a colored coin-to-virtual currency exchange did serve US customers (much like GLBSE attempted to in 2012 with bitcoin-only), and did derive income from fees on its exchange, they might need to register as an MSB and begin submitting SARs.[230]

This registration requirement nullifies one of the advertised advantages watermarked tokens had in the first place: the ability for non-accredited individuals to pseudonymously exchange virtual bearer assets without any gating process. It appears that based on the model above, to comply with regulations, many venture capital-backed watermarked token startups are effectively building permissioned-on-permissionless platforms.

(Schroeder 2015) also highlights this as it relates to UCC Article 8. Recall that one solution to "perfecting" bitcoin within the bounds of UCC statues is for a securities intermediary to take possession of a virtual currency. But in doing so, this removes one of the key purported benefits of a permissionless blockchain:

> If the owner of bitcoin were to choose to hold it indirectly through a financial intermediary, then she and the intermediary could elect to have it treated as a "financial asset" which is super-negotiable. Unfortunately, this comes at the cost of eliminating one of the primary attractions of cryptocurrency, namely the ability to engage in financial transactions directly without a third-party intermediary.[231]

If this is how the regulatory framework will continue to operate then organizations looking to use some type of distributed ledger or blockchain will likely see diminished utility using cryptocurrency-based systems and may gravitate towards permissioned networks that were designed with known and identified participants from day one.

Section 3.5: The Securities and Exchange Commission

On August 30, 2013, Mary Jo White, chairwoman of the SEC, explained to the Senate Committee on Homeland Security and Government Affairs:

> Whether a virtual currency is a security under the federal securities laws, and therefore subject to our regulation, is dependent on the particular facts and circumstances at issue. Regardless of whether an underlying virtual currency is itself a security, interests issued by entities owning virtual currencies or providing returns based on assets such as virtual currencies likely would be securities and therefore subject to our regulation.[232]

To date, the SEC has only brought action in a handful of virtual currency related cases, including Bitcoin Savings and Trust, SatoshiDice and most recently with Gemcoin, HashingSpace and GAW Miners.[233]

While nobody appears to be claiming that bitcoin itself is a security, it is still possible to run afoul of laws requiring registration of assets if they meet the criteria of what is defined as a security. This includes stocks, bonds, options, certain certificates of deposit, investment contracts, transferable shares and many more.[234]

There is a laundry list of items and the most litigated currently involves the characterization of an "investment contract" typically through the *Howey* test.[235]

Recall that an investment contract under the *Howey* decision involves:

1) the investment of money

2) in a common enterprise
3) with the expectation of profits
4) from the efforts of others.[236]

While a number of other Bitcoin-related endeavors, such as Sand Hill Exchange, have received paper work and settled with the SEC, none of the "Blockchain 2.0" projects that have raised funding through "crowdsales" have.[237] This is notable because many of the projects, such as OmniLayer, Swarm, as well as others which utilize the Bitcoin network, also rely on the appreciation of value of the underlying token in order to build out and fund their platforms.[238] This would seem to fulfill the "investment contract" characteristics of the Howey test.[239]

What if a crowdsale explicitly includes software within a pre-purchase agreement? Then it may not meet the 'investment standard' as defined by *Howey* however this has not been tested in court.[240]

With the advent of Kickstarter, Indiegogo and with the adoption of crowdfunding rules by the SEC in late October 2015, prefunding projects by pre-selling tokens is still a gray area.[241] Whether it is explicit or not, it is illegal if an individual issues an unregistered security via colored coins or anything else as it is the act itself that is unlawful.

What if a public company used a system such as Counterparty to issue and track registered securities?

This ties back into the discussion earlier in the paper. From a technical standpoint it can probably be done, but from a practical standpoint it runs into five practical challenges:

1) By requiring identification of all parties in such a transaction, these individuals and organizations lose the benefit of pseudonymous consensus that a public blockchain creates.
2) Public blockchains are currently significantly limited in the capacity and throughput. If even just one large securities exchange in the US attempted to hash all transactions of each trade onto the Bitcoin blockchain, there is no quick way to upgrade to handle the incoming volume and the network would grind to a halt as described in Appendix A.
3) There is no clear terms of service for using the Bitcoin network let alone a service-level agreement with mining farms/pools. In the event that a TOS or SLA is created, it would likely require identifying the parties which again, removes the utility that comes with pseudonymous interactions (i.e., users would be paying for Sybil protection when there are no longer Sybils as the network became gated and trusted).
4) As described in (Teo 2015), miners could be incentivized to "hold-up" the network, seeking higher rents to protect assets whose value is disproportional to the underlying carrier value itself which could lead to a log-jammed network.[242]
5) Due to the tragedy of the commons, it is unclear if the maintenance of public blockchains are sustainable in the long run. If they are not, then the issuance, tracking

> and trading activity – even if legal – that took place in the past via platforms on top may end up residing on an insecure or dead blockchain.

Whether or not these startups and projects involving initial coin offerings (ICOs) operate as an unregistered broker/dealer is also topic beyond the scope of this paper.

Section 4: Conclusion

Over the past 18 months, there has been an increased discussion about the utility of blockchains, distributed ledger technology, decentralized cryptocurrencies and shared ledgers. In fact, to the dismay of grammarians and empiricists, the term "blockchain" is now frequently used as a non-articled noun, existing on its own linguistic plane.

Consequently, the media is asking a series of questions. What can be put on "blockchain?" What can "blockchain" do for a specific industry? How can a company create "blockchain" strategy? What will the world look like when all identities, transactions, communications, documents and media are on "blockchain?"

The *reductio ad absurdum* is such that everything, including collectible card games, are being lofted onto "blockchain."[243] Full blockchain ahead, never mind the scalability, security, economics or legality of such a plan.

It may be the case that "decentralization for decentralization's sake" or "consensus for consensus's sake" or "blockchain for blockchain's sake" is a misguided approach to reorganizing bilateral and multilateral interactions. In some cases, there are other more appropriate solutions to a problem set.

Casey Kuhlman, co-founder of Eris Industries, shares his view:

> Distributed ledgers alone will have marginal impact. However smart contracts held on distributed ledgers will have a significant impact on our ability to provide assurances of process for distributed teams. This will be a game changer for companies big and small and for non-profit organizations big and small.[244]

Smart contracts were purposefully avoided in the conversation of this paper, despite the fact that the term is typically used in many conversations involving "blockchain" and watermarked tokens. Ironically, there is no consensus on the term or what it even legally represents, perhaps this will be rectified in future research.[245]

What, then, is the verdict?

Many of these watermarked projects appear to have arisen out of a sunken costs fallacy. Over a billion dollars has been invested in Bitcoin mining since 2009, and an additional $900 million in venture capital has gone into the greater "Bitcoin ecosystem" the past three years. It would be a shame to throw this away and start all over, right?

One manifestation of this fallacy includes hash-based proof-of-work inside financial institutions. One such scenario, proposed in (Bitfury 2015b, 2015c) is the following.[246]

Imagine a group of financial institutions or enterprises, such as banks, begin to create and use their own permission-based shared ledger. What happens if they are hacked and one of the members, Bob, goes rogue and begins to spam the network, or attempts to change the history of the ledger? Surely the only solution is to increase the difficulty rating and require all ledger members to use proof-of-work equipment?

Proof-of-work is unneeded because all members will know who is abusing the network, due to the fact that Bob signs all transactions with the same key(s). This is a type of man-in-the-middle attack to which Bitfury and others think permissioned networks are vulnerable, because they are assuming that the end user is not cryptographically signing the transaction. However, members signing transactions does not require running a full node, let alone a mining farm.

Thus, in this situation, white listing and black listing – as is already commonplace in financial networks – is effectively the same as increasing the difficulty rating. The membership network already knows who is who, and therefore proof-of-work is a vestigial, anachronistic option for permissioned systems.

While Bitcoin may appear safe enough to handle on-chain transactions with its own native currency (bitcoins), it is currently not a platform that could securely, economically or legally handle an *ad hoc* layer in the form of graffiti-tokens.

While many of the early adopters of blockchain-based technology believed that their anarchic protocols and networks would force courts to conform to their "code-is-law" point of view, the stark reality is that a ledger transferring off-chain titles and securities is only authoritative if courts recognize it as such.[247]

Public blockchains using watermarked platforms do not appear to meet this requirement. (Mizrahi 2015) echoes a similar conclusion regarding ownership changes and litigation: "A system which cannot address this issue can be impractical… A registry should comply with court's orders to reassign ownership."[248]

Or, as George Fogg explained several months ago:

> My libertarian friends have a belief they have created something that is outside of any statutory governance, and my response is you have created something novel that can help in transferring value across borders but you can't pretend that the UCC doesn't

> exist and because it does exist it affects bitcoin. Bitcoin is governed by the UCC. You can be an ostrich and pretend that it's not covered by it, or you can address that it is in fact covered by the statute and find a way to solve the problem.[249]

This candid view is probably one of the explanations underlying a recent trend in fintech startups: pivoting away from issuing assets directly onto public blockchains, and instead building their own permissioned ledgers.[250]

While it is debatable as to what direction(s) should be taken to securely scale decentralized cryptocurrencies, due to the aggregate issues in the previous sections, financial institutions and governmental organizations should not consider using these permissionless platforms for any financial activity that requires secure settlement finality of off-chain assets.

Part of this is due to the inability to map endogenous virtual outputs (watermarked tokens) onto the existing off-chain legal system, which is exogenous to the network and therefore cannot be legally or economically secured.[251] Due to the legal framework of KYC and AML, venture capital-funded startups invariably appear to recreate a permissioned system on top of an already expensive permissionless system.

Large organizations may be slow to adopt new technologies, especially ones that they do not understand. Bitcoin has demonstrated of the potential power of blockchains and distributed ledgers that is difficult for decision makers to ignore. It does not make them immediately practical, but it does make them interesting enough to explore and nurture.

And what about issuing watermarked tokens?

Based on its current incarnation, the Bitcoin network cannot provide the secure settlement finality for off-chain assets that many startups have recently proposed. Ignoring the challenges surrounding the tragedy of the commons governance, it makes little sense for financial institutions – which, for compliance reasons, need to document customers and partners – to build expensive cost drivers like a proof-of-work mining farm, yet gain no utility from pseudonymous consensus.

There is still room for both permissioned and permissionless blockchains (and whatever other term arises) to coexist, but that is beyond the scope of this paper.

Acknowledgements

I would like to thank the following individuals for their time, comments and feedback: Anton Bolotinsky, Arthur Breitman, Kathleen Breitman, Richard Brown, Vitalik Buterin, Dustin Byington, Preston Byrne, Christian Decker, Justin Dombrowski, Jacob Farber, L.M. Goodman, Ian Grigg, Peter Jensen-Haxel, Mike Hearn, Dave Hudson, Chris Huls, Casey Kuhlman, Jo Lang,

Jonathan Levin, Antony Lewis, Juan Llanos, Zaki Manian, Jared Marx, Todd McDonald, Ayoub Naciri, Piotr Piasecki, Raja Ramachandran, Robert Sams, Amor Sexton, Ryan Straus, Ernie Teo, Albert Veliz, and John Whelan.

Disclaimer

Tim Swanson's contribution to this publication was as a paid researcher for R3CEV, and was not part of his responsibilities with other organizations. He is reachable at tim@r3cev.com

Appendix A: Internal governance

Lack of explicit governance is reason alone to reconsider linking mission-critical off-chain financial products on top of public blockchains, which are economically and politically impacted by block sizes.

Because Bitcoin is a public good, there is no *de jure* entity to fire, hire or make decisions about how its blockchain should evolve (or not). While this makes it difficult for government actors to disrupt, it also creates uncertainty for developing and deploying a unified roadmap, insofar as creating new features or extensibility.

For instance, due to disagreements between Bitcoin Core developers, some new ideas and features get tested out and implemented in altcoins and alternative blockchains (e.g., ring signatures in CryptoNote, zk-SNARKs in Zerocash, group signatures in Tembusu).[252,253]

In practice, lack of clear governance and formalized structure devolves into an informal structure through factionalism and tribalism between special interest groups (some of whom appeal to ill-defined "social contracts" or *ex post facto* "constitutions" to justify their positions).[254,255]

In the case of Bitcoin, most of the discussion on governance and decision-making has degenerated into lobbying companies, organizations, and individuals via social media to promote one agenda over another.[256257] Recently, this has culminated in a power struggle between a handful of people who have committed access to the Bitcoin Core repo, in what is dubbed the "block size debate."[258]

What is the block size debate?

This challenge has been known for several years. In its current manifestation, Bitcoin Core (specifically, mining pools running Bitcoin Core) does not support block sizes larger than 1 megabyte (MB). The characteristic itself can be arbitrarily changed. Originally, there was no

limit. This was changed in 2010 from 33.5 megabytes to the current 1 MB due to spam concerns.[259]

What is the issue related to spam?

In one scenario, a farm that has invested capital in building an efficient network with its hashers could create an extra-large block and use it as a "weapon" by sending it to marginal pools and farms that must spend additional time unpacking it, which creates an opportunity cost.

Due to network propagation (Decker and Wattenhofer 2013; Levin 2014; Buterin 2014; Heilman et al. 2015; Garay et al. 2015), these smaller pools spend more time looking at older blocks than working on a new block, and thus their orphan rates measurably increase.[260] The possibility of 'selfish-mining' is also exacerbated by larger blocks, as noted in (Luu et al. 2015) and several others.[261]

There are economic and political trade-offs to increasing (and decreasing) the maximum block size allowed on a proof-of-work blockchain:[262]

- Keeping a 1 MB block size will likely require higher fees to on-chain end-users, but allegedly results in a topologically more decentralized network and less "spam" (e.g., fewer 'long-chains').
- With a larger 8 MB block size, such as the one used in the BIP 101 standard, mandatory fees to miners are lower for end-users. However as noted in (BitFury 2015a) this may come at the cost of fewer validating nodes on the network potentially reducing privacy.[263]

Over the past few months, several different block size special interest groups have coalesced, each with differing goals. The schism has also prompted worldwide media attention and created an expensive opportunity cost.

While not necessarily binary, there are two fears within the development community that (Brown 2015c) classifies as: "fear of practical failure" versus "fear of technical failure."[264]

For instance, one distinct developer group is working towards a 'hard fork' called Bitcoin XT, which implements the BIP 101 block size expansion plan.[265,266] BIP stands for Bitcoin Improvement Proposal.

Support behind this effort includes many Western VC-backed companies such as Coinbase that have publicly claimed that Bitcoin-based companies will one day be able to compete with payment incumbents (such as PayPal and Visa). Therefore, they would like to modify Bitcoin achieve a similar transaction-per-second metric.[267] To do so under the BIP 101 proposal, they would like to double the block size every two years, starting at 8 MB beginning in January 2016.

This segment of the community would likely be defined as trying to solve the first fear.

In their view, as blocks become consistently full (leading to longer wait times as transactions compete to be filled into scarce space) and fees begin to rise, casual users will leave the ecosystem for other platforms. In Brown's words, their fear is that "failing to increase the blocksize, a change which has uncertain risks in any case, will drive away users and make the system a failure in all practical cases."[268]

Over the past summer, interest in the Bitcoin XT project has markedly increased. One way to measure this is by looking at the amount of full nodes (approximately 10%) that now run Bitcoin XT, as well as mining pools that sign support for bigger blocks in the coinbase transaction.[269,270,271]

Another group are promoting a more conservative approach to increasing the blocksize.[272] Support behind their vision includes a variety of companies such as Blockstream, and they are typically more inclined to support either a low increase or none for the time being.

This group would likely be defined as trying to solve the second fear. According to (Brown 2015c): "the change, which has uncertain benefits, could cause catastrophic harm. Better not do it just yet."

Some of these same developers prefer an alternative method of scaling by attempting to build a proposed "Lightning Network" (based on payment channels), as well as off-loading a lot of the transaction volume to sidechains. Blockstream is building a few varieties of this (e.g., federated peg and two-way peg).[273]

There are several other contingents with unique viewpoints, including a large portion of the Chinese mining community (which collectively represents about 55% of the network hashrate), each of whom have come out with differing opinions. For example, F2Pool, currently the largest mining pool, is outspoken in its opposition of Bitcoin XT, calling for its boycott.[274] In contrast, Slush's pool, a Czech-based technology company that also created the first public mining pool, is in favor of BIP 101.[275]

There are other also active, independent developers that have created alternative block size proposals, such as Meni Rosenfeld ("elastic cap"), Jeff Garzik (author of the BIP 100 plan) and Peter Todd (tree chains), each of whom has elaborated on the trade-offs that both larger block sizes and a hard fork will have on the Bitcoin network.[276]

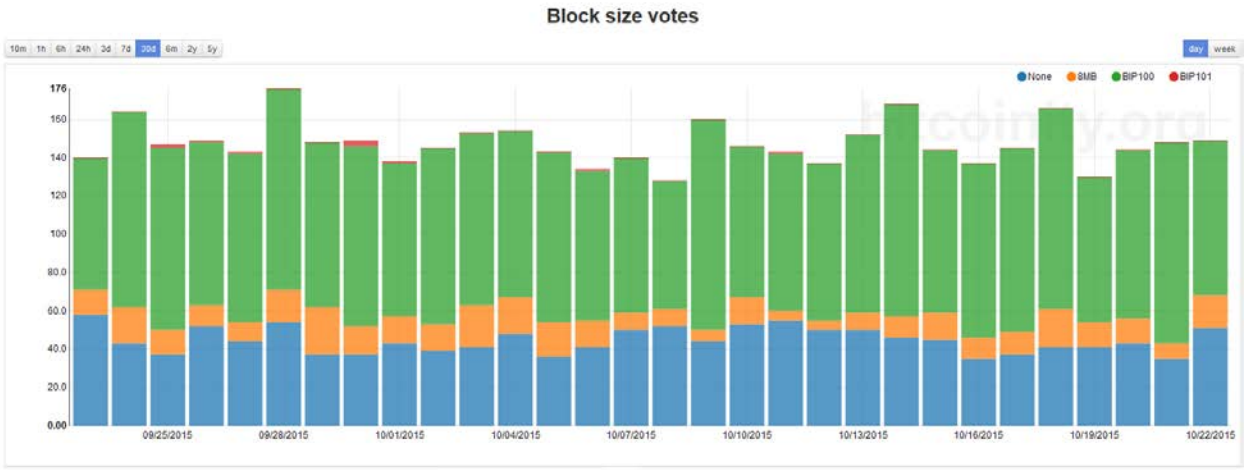

Block size vote. Source: Bitcoinity.org

Mining pools have the ability to sign a message in the coinbase transaction when they win a block reward. Some of them have used this space to voice their support of certain proposals. At the time of this writing, as measured by miner "votes," BIP 100 appears to be the leading proposal although it has not been built or implemented.[277,278]

What does this have to do with colored coins?

While it may sound like an academic debate, all of these political decisions impact and potentially disrupt *any* layer that relies on the Bitcoin mining network to secure it. Thus it directly impacts watermarked token platforms as well.

Why is a block size an issue at this time?

At current usage rates (shown in the chart below), blocks will be consistently filled by December 2016, and the network throughput is likely to max out at around 2.8 transactions per second.[279]

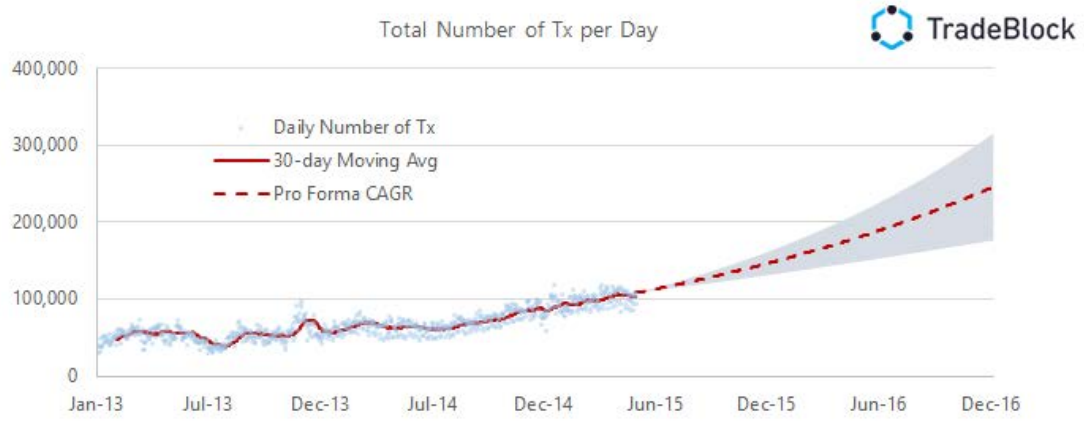

Source: TradeBlock.com

Even with a variety of proposed tweaks such as increasing block sizes, (Decker and Wattenhofer 2015) noted that:

> Bitcoin in its current form will have a hard time scaling beyond 100 transactions per second, because of storage, processing, latency, and bandwidth. The problem of Bitcoin is its reliance on a synchronized global state, the replicated blockchain.[280]

Thus, for the developers in the "fear of practical failure" group, not being able to scale beyond 100 transactions per second is a serious hurdle they aim to overcome. How do we know that this is an issue that should be looked at today?

In late May 2015, at least one actor (possibly CoinWallet) proceeded to send spam-like transactions onto the network to determine the economic and therefore the technological costs of larger blocks.[281]

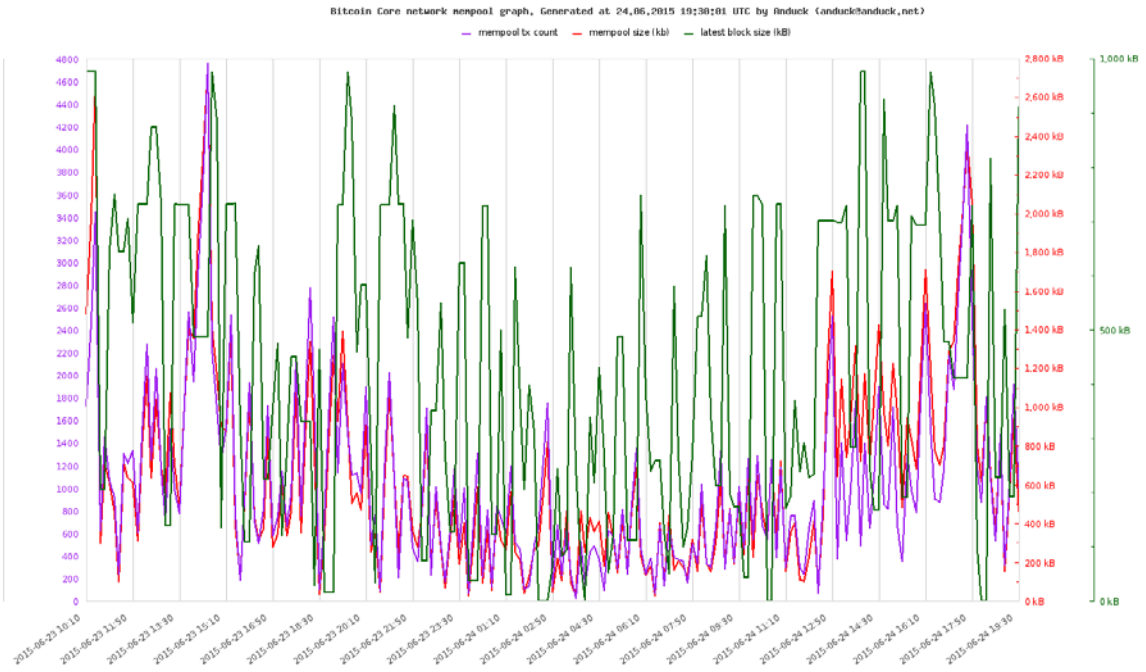

Source: CoinWallet.eu

During this test, thousands of 700 KB transactions were sent within the first 20 minutes, each one relatively small (0.0001 bitcoin). The mempool, where unconfirmed transactions reside prior to being included into a block, ballooned to more than 14,000 unconfirmed transactions. Consequently, a variety of APIs broke, including the Blockchain.info API which the team was using to conduct the test.[282]

What they found is that today, for less than 2 bitcoins (€434) in fees, an actor can disrupt and clog the network for hours. For some colored coin-based stocks, it would be a disaster; particular trades could stop suddenly due to similar attacks.[283] As of this writing in September

2015, CoinWallet is conducting a longer follow-on test with the aim of motivating people to implement larger block sizes.[284]

Beginning on July 6, 2015, another test (likely from a different actor) began sending a continuous stream of thousands of transactions.[285]

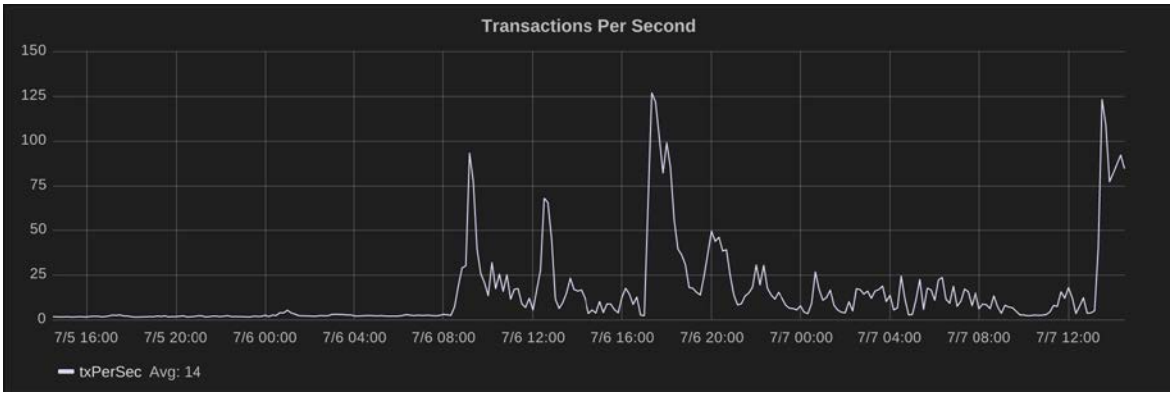

Measuring the second stress test in transactions per second. Source: Statoshi.info

At its peak, and while attempting to process 125 transactions per second, more than 100,000 unconfirmed transactions sat in the mempool; a backlog that took more than a day to fully plow through. According to one independent calculation based on how the fee market changed upward during this test, if an actor wanted to commit an on-chain transactional "attack" on the network, it would cost roughly $4,536 per day to continually fill up the 1 MB block size.[286]

But if fees increase over time, doesn't that reduce the inclusivity of Bitcoin?

Most of current blockchain traffic is effectively "spam" (e.g., "long chains") that would likely not exist with higher fees.[287]

But more to the point, and due to the relatively centralized network graph, the reliance on specific "super nodes" from API providers and mining pools (Heilman et al. 2015; Gervais et al. 2015; Miller et al. 2015) at the current trajectory, the various colored coin and metaplatforms attempting to latch onto Bitcoin will likely make them a victim of their own success – the more top-heavy the network becomes, the larger the incentive for attacking nodes, pools, API providers, routers and other infrastructure.[288,289,290]

The block size debate has also opened up the possibility of future governance changes via populism.

For instance, if developers were able to create a new 'hard fork' with a different block size due to increased interest on social media, what about changing the purported hard cap of 21 million bitcoins as explored in (Goodman 2014)?[291] In the event that a real fee market does not replace block rewards, would mining pools and large intermediaries be willing and able to fork the code once again, replacing the asymptotical money supply with one with permanent

inflation? How would this impact the fee and security structure of watermarked token platforms?

Since fees are not increasing as the popular narrative or "folklore" predicted, and miners are still heavily dependent on seigniorage to maintain security; how could governance be impacted by future purposeful attempts at forking the network to maintain the existing inflation rate?[292]

Was this bound to happen? The governance issue was independently explored several times in the past including by (Saberhagen 2013, Goodman 2014).[293]

Because technology and usage are not static, there must be a way to clearly upgrade and update both the software and network. Currently, the BIP process is an *ad hoc* kludge that depends on altruism and charity, neither of which are sustainable– and, as shown empirically, are beholden to special interest groups and their stakeholders.

The current situation has been likened to the Cold War missile crisis, where brinkmanship between groups claiming to represent both large holders of bitcoin and hashrate has resulted in détente.[294,295]

And what are the consequences of détente?

According to L.M. Goodman, creator of the Tezos financial platform:

> When it comes to long term survival, adaptability is more important than strength. Seeing distributed ledger as mere technology is shortsighted, they are first and foremost networks and, as such, their governance model is paramount to their success. A decentralized network that does not internalize its governance is condemned to stagnation or centralization.[296]

Other networks have learned from Bitcoin's governance limitations.

For instance, Peernova is building its platform that includes version control and Tezos has created a self-amending chain so that governance is endogenous. Similarly, Ethereum, while itself a public network, has created a formal power structure (a foundation incorporated in Switzerland) and has explicitly laid out its hierarchy, as well as a general roadmap of the future (e.g., switch from proof-of-work to proof-of-stake; private organizations being tasked to take leading roles).[297,298,299]

De facto governance

It is unclear at this time what fork or forks will arise and become the new "consensus" chain. Because they rely on Bitcoin miners for security, watermarked token platforms are directly impacted by such decisions.

To prevent a worst-case scenario in which the network is either attacked or a large minority successfully creates a competing chain (e.g., one with a vibrant alternative ecosystem), the

community began hosting a series of “scalability workshops” inaugurated in September 2015 in order to build a consensus around code, reviews, and testing.[300]

Others have proposed forming a new non-profit board, or “voting,” through movement of bitcoins or even social media polling, although these are problematic due to self-promotion attacks.[301]

One such non-profit board proposal (Mougayar 2015) asks:

> Why not fund a core team of people, including developers and others, and roll them into an independent non-profit organization that has credibility and neutrality. Why not skim 1% of the $400-600M transaction fees and fund a stable organization?[302]

The fees in this case equal the aggregate market value of the block rewards plus transaction fees (which are still negligible, roughly 0.5% of all revenue a miner receives).

Ignoring the coordination and collective action problems of convincing miners to redistribute their income, there may be some legal issues to this proposal. For a decentralized cryptocurrency which aims to be censorship resistant, creating an ICANN, IEEE, IETF, or other formal standardization body could create as many problems as it solves.[303]

If, as (Mougayar 2015) hypothesizes, a new centralized organization is created with identifiable actors that design and administer this code, then it may meet the FinCEN definition of administrator.[304,305]

What about other intrinsic (on-chain) versus extrinsic (off-chain) solutions to this governance impasse?

For instance, on August 17, 2015, Brian Armstrong, CEO of Coinbase, tweeted that:

> A lot of people might misunderstand a bitcoin "fork" as meaning two versions of bitcoin would exist indefinitely, which would be bad. [T]he reality is more like an election, there is one clear winner, decided likely within a few hours, in bitcoin you vote on upgrades. [T]he longest chain is by def bitcoin, with a few months warning for when the election happens, users would not experience disruption. [S]ervice providers like Coinbase would have time to prepare for either outcome. [H]ard forks probably shouldn't happen frequently, but periodically they are an elegant solution that helps bitcoin keep growing. [I]t's a huge advantage that bitcoin has a built in voting process to handle upgrades, checks and balances on maintainers of any fork, etc. [306]

While miners and nodes can and do “vote” on which chain they think should be the canonical one, there is no explicit, endogenous governance structure within Bitcoin itself. Hard forks are not *ipso facto* elegant solutions. The current “missile crisis” facing Bitcoin is arguably not an advantage as lack of clear, definitive governance creates uncertainty in the marketplace.

There are proposed methods that developers on the edges can use to concoct "social" interactions with the Bitcoin network, but this is *ad hoc* – not by design – and fully dependent on solving exogenous coordination problems.

What about voting via social media and forums?

What is a self-promotion attack? (Hoffman et al. 2007) states that attackers can manipulate their own reputation by falsely increasing it. [307] This is a problem that plagues social media networks such as Reddit and Twitter, but also older systems like the Maze file sharing system.[308]

(Haynes 2015) specifically looks at such type of manipulation on Reddit, where much of the block size and Bitcoin governance debate was, and is, taking place.[309]

The author used machine learning and authorship attribution techniques "to determine whether comments from user accounts that are active in the debate are from the same author. The techniques used are able to recall over 90% of all instances of multiple account use and achieve up to 72% for the true positive rate."[310]

In other words, individuals and organizations are using "sock puppets" on social media to manipulate the Bitcoin debate towards specific viewpoints.

Is this solely related to Bitcoin? No. In July 2007, it was discovered that John Mackey, co-founder of Whole Foods, had used at least one sock puppet on a Yahoo forum to talk up his company and talk down the competition, Wild Oats Market.[311] Thus, research into authorship analysis extending into the Bitcoin startup ecosystem itself – and its enthusiastic supporters and founders – is likely a ripe area for future scrutiny.

One tradeoff with building a public network that is difficult to censor is the evolution of an unpredictable governance apparatus, which fills the decision-making vacuum.

A financial network is different than an information network. It appears that the cryptocurrency community is getting a crash course in the challenges that arise when bootstrapping a financial network through internet forums and internet relay chat (IRC), without fully quantifying or qualifying the organizational issues traditional financial service firms must be cognizant of as they face them on a daily basis.

And any platform – such as a colored coin system – that resides on top of a public blockchain lacking clear governance will be impacted by these challenges.

Appendix B: A straw man game theory model

Bitcoin's security assumption is based on the cost of gaining control of a majority of the network hashrate. For instance, what would be the cost of reversing a $1 million transaction represented by a watermarked token? Would it also be $1 million?[312]

In the original white paper, Satoshi deemed 6 confirmations as being "secure." That is to say, if a transaction was buried under a sufficient dollar amount of proof-of-work, it is supposed to be considered irreversible.[313]

What do we mean by the cost of reversing a transaction?

To successfully disrupt the network, the *maximum* cost to do so is roughly 0.5 x MC, where MC is the marginal cost of production.

Is there a way to visualize this?

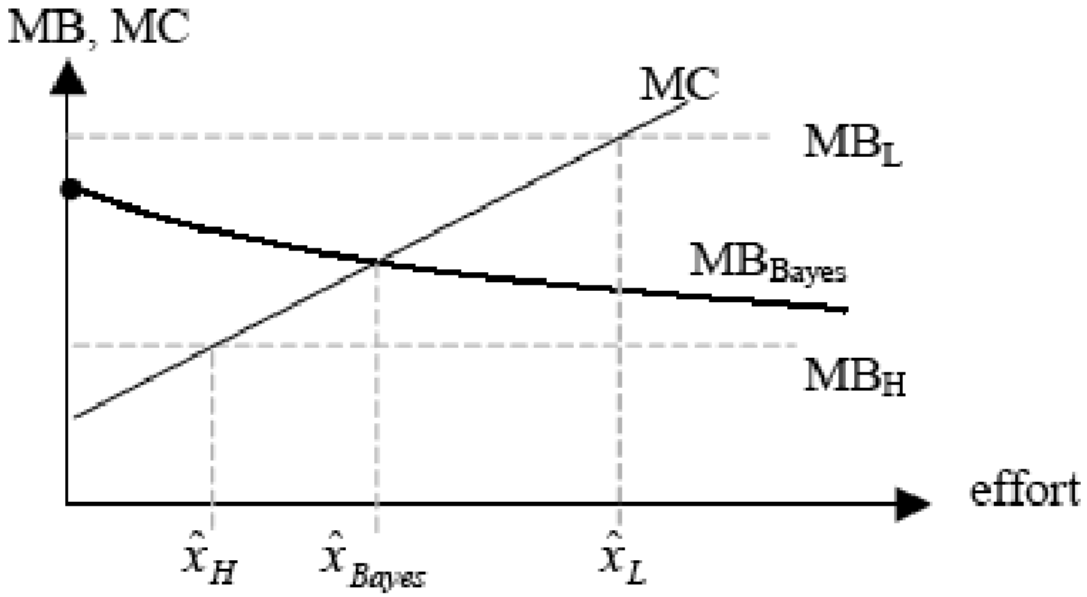

Source: Cremonini and Nizovtsev (2006)

While not originally intended to analyze attacks on cryptocurrency networks, the diagram above visualizes the tradeoffs attackers have, between their marginal benefits of attack and marginal costs. The generalized model attempts to explain the opportunity costs an attacker on a network has, or as (Cremonini and Nizovtsev 2006) explain:

> Marginal benefit and the solution to the attacker's optimal stopping problem for the complete information and the incomplete asymmetric information cases. The dot represents the prior probability of success from a randomly selected target. The

> expected marginal benefit of attacker's effort decreases because the more effort the attacker spends on a target with no success, the more he believes he is dealing with an H-type system.[314]

An H-type system is a target system with a high security level (in contrast to a low security level). In their model, high and low security systems have different incentives for transparency and opacity.

It is this point that becomes relevant to the discussion of cryptocurrencies. If the MB relative to the MC of an initial attack changes by several orders of magnitude, as in the case of watermarked tokens, then it may be worth their time and resources to actually attack the network itself instead of the edges (as historically has been the case).

In the author's words:

> If switching costs are a representation of the amount of effort an attacker needs to spend to find out the target type, then the above result means that the efficacy of security investment depends on the characteristic of the environment which we call "opacity". The greater the switching costs, the more "opaque" the environment is in the sense that it gets harder for attackers to determine the type of a target, which weakens the behavioral effect. If, on the contrary, the environment is "transparent" and determining the type is relatively easy, then the behavioral and the overall effects of a security investment will be stronger.

Empirically, we know there is a spectrum in which high security can become low security overnight, given a sufficiently good zero-day attack.[315]

If colored coins and metacoins are widely adopted and use the built-in scripting methods, then there is potential for a seemingly unlimited amount of "assets" to be traded on the Bitcoin network. If these several thousand colored coin users add additional value, this creates an incentive for attackers to attack the network through colored coin-based double-spending attacks. Whether or not the assets themselves are transparent or opaque is briefly explored later.

It bears mentioning that this was another objection identified in (Rosenfeld 2012). The author wrote that:

> More difficult than the problem of funding network resources, is the problem of funding hashing. Since the cost of hashing is amortized over all transactions, it is essentially a bargaining game between miners and users, which unconstrained would lead to a race to the bottom. As such it will be useful to have protocol-enforced methods (such as a limit on the total value of transfer per block) to make sure fees are paid by those who can afford them (typically senders of high-value transactions). With colored coins, the network is unable to determine the value of the transactions and charge accordingly.

> This means what this economic activity contributes to the hashing network is not proportional to, and typically lower than the value it obtains from it. However, no actual harm will be done by colored coins; and while a system which does contribute fairly would be preferable, unless such a system can come up with, this objection is moot.[316]

Because it is a public good, the proposal above runs into another coordination and collective action problem. Who gets to decide what the limit on total value of transfers is? What is the governance structure around how this is decided?

Since this is a positive claim, the onus is on those marketing watermarked solutions as secure. Or as Bruce Schneier, an information security expert, would say: assume all systems are insecure until proven otherwise.[317]

To-date, there have been no publicly released studies on the feasibility or security of watermarked systems, yet there are more than a dozen startups attempting to build financial platforms on an untested sand castle.[318] As (Garaye et al. 2015) state, "[…] a thorough analysis establishing the exact security properties of the Bitcoin system has yet to appear."[319]

Christian Decker, a distributed computing expert at ETH Zurich and co-author of the only research paper (Decker et al. 2014) to broach the topic of security and colored coins, said the following:

> We believe that some use-cases that are currently piggybacked on top of Bitcoin do not participate in incentivizing participation. Our protocol would allow experimenting with these kinds of alternative applications in their own context, keeping data separate and without polluting the "main network". Not only is this more secure but also cheaper than attempting to build their own blockchain, which we see as the main reasons for other applications to mix with the Bitcoin blockchain.
>
> We have not seen any in-depth analysis on the impact of multiple applications mixed into the main Bitcoin blockchain, but it would certainly be interesting to see.
>
> We definitely need a clean separation of applications, without attempting to hide one application's information in another, that will allow us to evaluate the incentive structure necessary for each application individually.[320]

What is one way an attack could be done?

A rogue attacker could sell an asset and build a competing tree (consensus in Bitcoin is based on whatever is the longest tree of blocks).[321] After a successful 51% attack, the rogue attacker could then broadcast a fake chain built without the corresponding asset, having switched it out and thus effectively double-spending. And if the total value of the network's transactions has

at least twice the value of bitcoin, then there is a financial incentive for rogue participants to attack the network.

This scenario raises the question: what, then, is the potential divergence in value between bitcoin the currency and Bitcoin the network (which can transfer and protect other data)? This issue only presents itself now, as previously, only bitcoins – and no other apps, assets or instruments – existed on the network.

This also gives rise to a collective action and coordination problem, because miners have to keep track of the color, the exchanges the color is being traded on, and the settlement price (if there is one), so they can adequately gauge market clearing prices and readjust the coinbase reward every 10 minutes.

Even if this coordination problem is solved, the seigniorage reward itself does not dynamically adjust – the current fixed income does not reflect the actual value being transacted on the network. So colored coins on a fully decentralized network could end up on an *undersecured* network of their own making with the only solution: recode the block rewards based on the value of the color. This presents a number of technical and social engineering challenges.[322]

And what happens if there is a block reorg?

That is to say, due to its nominally decentralized topology, and to paraphrase what Henry Kissinger allegedly once said, "Who do I call if I want to speak to Bitcoin?"[323] What is the Bitcoin network's service-level agreement?

In other words, while these new financial instruments could technically be exchanged to other participants, the current network cannot automatically incentivize their protection or account for their social value, the equivalent of using a mall security guard to protect Fort Knox.

While miners may be able to protect against amateurish shoplifters or even unorganized cat burglars, once organized criminals calculate and realize that one "color" asset is worth the economic effort of attacking the vault, they may try to do so.

This vulnerability is not new for P2P networks.

In (Hoffman et al., 2007) the authors looked at attacks on reputation systems and peer-to-peer applications to build a classification model. They note that:

> We assume that attackers are motivated either by selfish or malicious intent. Selfish (or rational) attackers manipulate reputation value for their own benefit, while malicious attackers attempt to degrade the reputations of others or impact the availability of the reputation system itself.

> In general, attackers can either work alone or in coalitions. Although both scenarios are possible and relevant with respect to reputation systems, we are primarily concerned with attacked caused by coalitions of possibly coordinated attackers. Coordinated attacks are more difficult to detect and defend against because attackers can exhibit multi-faceted behavior that allows them to partially hide within their malicious coalition. We note that while coordinated attacks caused by multiple Sybil identities are a danger to any distributed system, even without such attacks, reputation systems are vulnerable to malicious collectives of nodes.
>
> We consider all attacks to be active since any form of attack on the reputation system requires interaction with the system, such as injecting false information, modifying entrusted information, refusing to forward information, deviating from the algorithmic processes, or actively attempting to subvert the availability of the system.[324]

Basically, the challenges facing many older peer-to-peer, reputation-based decentralized networks are the same challenges that Bitcoin faces. And coordinated attacks, including organized crime, cannot be ruled out.

That is an implicit negative for investors and users, and raises some concerns in the future. If a party has the ability to invalidate Bitcoin accounts based on their own criteria, the miners might gain an influence over the colored coins, and may bias various aspects of the economy incentivized through some kind of backchannel payment.

On-chain protection and control of off-chain assets

Colored coin issuance appear to respond differently to the game theory assumptions around which a seemingly homogenous cryptocurrency network, such as Bitcoin, was built.

The economic term that describes one set of potential vulnerabilities (discussed above) for watermarked token platforms is the "hold-up problem."[325] It bears mentioning that this is just one model that explains how interactions could potentially occur.

The hold-up problem is usually only a problem for systems which are *de jure* decentralized, but *de facto* ones have an extreme concentration of consensus power (such as mining). Such systems are likely ill-fated no matter what, due to the governance issues described in Appendix A.

The hold-up problem arises in situations where it is difficult (or impossible) to design complete contracts between buyers and sellers. That is to say, not all contingencies can be accounted for. Therefore, the contracts are incomplete, and one party gains leverage over the other due to some unforeseen change (e.g., shift in global supply chain due to a natural disaster or political upheaval).

This occurs in many industries globally.

An example might be if Alice, a car manufacturer, exclusively sources her steering wheels from just one supplier, Bob's automotive center. Knowing that Alice is unable to break the contract and quickly source to another steering wheel vendor, Bob uses his position as leverage to renegotiate better terms. In some cases, this type of scenario leads to underinvestment or lower quality because manufacturers are unwilling to take the risks of entering into such asymmetrical relationships. This is one reason why phone manufacturers such as Apple typically prefer to source parts from at least two vendors, so they are not put into a levered position.[326]

Watermarked token platforms on public networks have the same fundamental problem, due to the limited amount of suppliers (miners) that could have an incentive to censor transactions, collude, or even form a cartel as block rewards decline over time.

For instance, with colored coins, there is specific information that miners (mining pools) can use to hold-up transactions. That is to say, "color aware" miners can charge higher fees which, in theory, would lower their incentive to short the market. In order to prevent attacks, pools would need to receive information about specific colors.[327] This would lead to rent seeking from miners: the rents they extract would include the cost of the attack. This would, in turn, make fees significantly higher.

And when there is no information, or where pools are "color blind," then they cannot hold-up. While miners may not initially charge higher fees, they do know a certain percentage of bitcoins are, for instance, Apple shares, so they hold the network hostage.

In practice, cartels are hard to maintain without someone cheating (e.g., high prices lead OPEC members to cheat), or, in the case of cryptocurrency networks, breaking ranks and mining the colored transaction anyways.[328] The question as to which is more likely to happen – block reorganization or censorship of a watermarked token – is an open research problem.

If miners can never identify colored coins, then there won't be a hold-up problem at all. If miners *can* identify colored coins, yet if there are miners that do not care, then it is not a problem either.

The specific case (Teo 2015) explored is: if the value of the colored coins is sufficiently high, assuming miners are rational, they will want to extract the utility of it.[329]

Even if a colluding set of miners cheats, driving the colored coin fee lower, there would likely still be a non-zero fee.

(Teo 2015) presented one possible set of conditions and outcomes of a hold-up model.[330] The model does not cover everything, it simply illustrates some of the potential areas that could happen in the mining process. It demonstrated that there could be various strategic implications for colored coins, such as the issuer possibly having to become a miner, possible implications about transaction fees and possibility of attacks. For example, if Alice knows that

none of the existing miners want to mine Alice's transactions, it would make sense to integrate (become a miner).

The first step is for a colored coin issuer, Bob, to make a choice: whether or not to become his own miner. If Bob wants to be his own miner, he would need to be sufficiently large in order to get his transactions included. But that would require heavy investment in a mining farm. There is no free lunch.

And as explored in (Teo 2015), if the marginal cost of attacking the network is very expensive, it may outweigh the marginal benefit, and therefore no attack occurs. With this (and given the various scenarios) we can say under certain conditions that the colored coin processor will not become a miner, and the system will be attacked.

There are at least three limitations not explored in this model:

1) There is the question of whether the color coin issuer (CCI) will want to orchestrate an attack on itself. Although there could be an incentive to attack itself (e.g., "shorting" in the process but this may have legal implications), the one-off payment may be worth it to the miner.
2) It also assumes that miners can *quickly* and *efficiently* distinguish which bitcoins are colored or not; these watermarked tokens can be hidden in "real" transactions and it may be too time-consuming to identify.
3) Game theory itself rests on the assumption of *rational* behavior, which cannot always be taken as a given. Similar assumptions have been made to support various financial models, which have been proven to be shaky by behavioral economics (e.g., Kahneman and Tversky's Prospect Theory).[331]

Once Alice, a mining pool, attacks, there is no chance that she will be able to mine colored coins in the future, as the social impact would cause reputational risk (i.e., people would no longer use her pool).

Based on the current governance vacuum in most cryptocurrency systems, it is likely that there could be larger more important governance issues long before any type of hold-up scenario will occur.

Endnotes

[1] I would like to thank Robert Sams for clearly articulating the point related to off-chain titles.

[2] One area that is not discussed is the impact of watermarked tokens with respect to simplified payment verification (SPV). The Ethereum white paper discusses this limitation (p. 7-8), however for financial institutions, finding a trusted node to give you a snapshot of the ledger probably is not too problematic. Because, in practice, we are talking about large financial applications, not something that needs to run on a cellphone. Thanks to Arthur Breitman for pointing this out.

[3] Bitcoin was not the first cryptocurrency. See A Quick History of Cryptocurrencies BBTC — Before Bitcoin by Ken Griffith. The term token (and tokenization) is used differently in the cryptocurrency industry than it is in the data security industry and often causes confusion.
[4] Not all distributed ledgers are blockchains and vice-versa.
[5] Overview of Colored Coins by Meni Rosenfeld.
[6] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder.
[7] For an explanation of how some colored coin encoding works, such as Enhanced Padded Order-Based Coloring, see EPOBC simple and Iddo Bentov (2015).
[8] Both CoinPrism and CoinSciences have also announced non-Bitcoin, non-public distributed ledgers targeted at enterprises called Openchain and Multichain, respectively.
[9] Several of these colored coin startups have also had to pivot away from their original idea: operating a web-based exchange that allowed users to buy, sell and trade unregistered securities. They eventually moved away from that idea due to regulatory concerns.
[10] Why only seemingly homogeneous? The question as to whether or not all bitcoins are the same is a controversial. Because of taint analysis, it may be the case that certain discrete amounts of bitcoin are treated differently than others. Similarly, if certain discrete amounts of bitcoin are "colored" than may be treated differently by mining pools. Therefore in practice it may be more accurate to use the term *heterogeneous* units.
[11] The debate over whether or not all satoshis are fungible has been ongoing for multiple years. Currently, under US law, bitcoins are not deemed fungible by any governmental body. Perhaps in the future this will change. See also the last chapter in *The Law of Bitcoin*.
[12] Decentralised Digital Asset Registers – Concepts by Richard Brown.
[13] One reviewer pointed out a potential exception: the right of free individuals to make any (legal) contracts between themselves that they choose. For instance, the DMV might not recognize the transaction, but what's to stop a new corporate entity interposing itself such that it appears as the owner of thousands of cars on the DMV's database. Users of a blockchain system actually have a contract with that entity (i.e. they buy the car from that entity and the entity records the fact on a blockchain). If all participants in that scheme agree amongst themselves that this particular blockchain is the authoritative record for who owns which claims against that firm, then it may indeed be considered an authoritative ledger. See also p. 34 in Bitcoin and the Uniform Commercial Code by Jeanne Schroeder.
[14] The twist to the current batch of startups and entrepreneurial activity is that they are trying to move a securities entitlement by moving a decentralized token. Bitcoin creates a twist because it exists outside an account based system. In that way it is like dollars vs. checks, or certificates vs. book entry. From a legal standpoint in the United States, for it to be negotiable, the token must be issued by a "securities intermediary" under article 8 of the Uniform Commercial Code (UCC). As Ryan Straus, principal attorney at Riddell Williams, notes: "Property systems are information systems that connect specific people with specific things. Like other information systems, property systems incorporate design decisions about what to remember and what to forget. Unlike many information systems, however, the design decisions made with respect to property systems are enshrined in law, not code. Issuing a security and establishing a security entitlement is not the same thing, and a security intermediary seems to be required for the latter but not the former (i.e., securities entitlement are just claims against securities intermediaries). The underlying asset is held by the intermediary and credited to a securities account." I would like to thank James Duchenne and Ryan Straus for their insights on this. Personal correspondence, August 27, 2015.
[15] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder, p. 53.
[16] 80 bytes OP_RETURN explained from CoinPrism.
[17] Parasitic in the sense that these new platforms consume more system resources than provide in return. It appears that the first time the word "parasitic" was publicly used to describe this economic interaction was "Disentangling Crypto-Coin Mining: Timestamping, Proof-of-Publication, and Validation" by Peter Todd in November 2013.
[18] In September 2013 Mike Hearn also pointed out problems with the Omni/Mastercoin approach: "But I'm really unclear why you need to build things into Bitcoin anyway. If you or the people you're working with aren't able to make a merge mined coin with your current abilities, then maybe it's time to learn how?" and "I suspect the real reason you want to use normal Bitcoin transactions is you don't have the technical skill required to implement your ideas in a separate system."

[19] According to one developer, the Counterparty-embedded consensus system is distributed between 100-300 nodes (these nodes are run by independent groups such as: CoinDaddy, CounterParty itself, ShapeShift, Tokenly, Vennd). Dogeparty also exists and is a fork of Counterparty on top of the Dogecoin network**.** See Setting up a Counterparty Federated Node.

[20] On p. 7 in the Ethereum whitepaper, the author notes that one *purported* advantage for metacoins that live on top of Bitcoin, "This provides an easy mechanism for creating an arbitrary cryptocurrency protocol, potentially with advanced features that cannot be implemented inside of Bitcoin itself, but with a very low development cost, since the complexities of mining and networking are already handled by the Bitcoin protocol."

[21] Developers in both communities have created special software clients (wallets) that are color- or meta-aware, and that can be used to create and track these watermarked tokens. The key difference is that colored coins schemes, through the watermarking metaphor, attempt to create an equivalence between "bitcoin ownership transfer" and "colored asset ownership transfer." Whereas the ECSs simply use Bitcoin as a storage/proof-of-publication and that the ECS rules/protocol can be entirely arbitrarily interpreted through code running by the relevant ECS participants and opaque the underlying network. Thanks to Richard Brown and Anton Bolotinsky for clarifying this.

[22] A blockchain-based property ownership recording system by Alex Mizrahi (2015).

[23] Or dogecoins, in the case of Dogeparty.

[24] I would like to thank Arthur Breitman for clearly articulating this.

[25] Another potential option is via a separate protocol using some kind of pegged "sidechain." Though not necessarily the "two-way sidechains" being developed by Blockstream. In his first interview discussing two-way pegged sidechains that Blockstream is building, in March 2014, Adam Back, co-founder of Blockstream and creator of Hashcash explained one of the drawbacks of "colored coins," see: Paraphrased notes from Back and Hill interview.

[26] While Blockstream has announced (Liquid) but not released their "federated peg," one legal question is that if they do create a system of "trusted" validators, then whichever validator actually signs off or packages the transaction (or assets) could be serving as custodians of those assets. Stellar and Ripple Labs both may have a similar legal question with the validators they operate. See Blockstream to Launch First Sidechain for Bitcoin Exchanges

[27] While it may seem trivial to add a few hundred lines of code to enable asset issuance and other "coloring" features to layer 3 in Bitcoin, the coordination problem of gaining the approval or consensus from developers, as well as the governance thereof, has been shown to be the reason layer 4 platforms were created in the first place. And while it is possible to do so in a "private Bitcoin fork," the question is, why? If Bob's goal is to build an industrial-scale, closed-loop network, Bob will probably have to start from scratch with different security assumptions.

[28] The Ethereum white paper (p. 9) also singles out this scalability issue: "Thus, we see three approaches to building advanced applications on top of cryptocurrency: building a new blockchain, using scripting on top of Bitcoin, and building a meta-protocol on top of Bitcoin. Building a new blockchain allows for unlimited freedom in building a feature set, but at the cost of development time, bootstrapping effort and security. Using scripting is easy to implement and standardize, but is very limited in its capabilities, and meta-protocols, while easy, suffer from faults in scalability."

[29] *Prima facie*, there is nothing inherently wrong with a private Bitcoin-derived blockchain, but why are they using colored coins? If Alice is no longer using Bitcoin's mining network, why not build a new system that is based around the security assumptions and requirements of the new model? I would like to thank Arthur Breitman for this insight.

[30] Ethereum raised 31,529 BTC (exodus address) between July and September 2014, which, at the time, was worth around $18 million. See Launching the Ether Sale by Vitalik Buterin

[31] Because it was built from scratch, Ethereum is technically not a fork of Bitcoin. Also, it bears repeating that there is only so much you can encode by moving a small sum of money (a "colored coin") from one place to another: it is inefficient in terms of space taken on the block, and it is hard to attach data properly. A developer, Bob, does not have the automatic triggers that he will theoretically have in Ethereum where transactions can automatically trigger more transactions – Bob has to rely on workflow automation to do that which could makes auditabilit harder to accomplish.

[32] The inability to "enforce" external contracts is not limited to Bitcoin. Ripple's network likewise cannot "enforce IOU transactions." See p. 16 in Ripple: Overview and Outlook by Armknecht et al.
[33] One reviewer pointed out that it is very possible that some of the non-Bitcoin-based distributed ledger startups will build applications where the underlying protocol does not know what is in a transaction and yet can still validate it. One way that problem could be solved on a permissioned ledger is that the validating node(s) will maintain a separate database that ties a public address to a customer account. While a Bitcoin-related company could do that as well with a public blockchain, it would remove the utility surrounding pseudonymous interactions. Arguably it would be self-defeating for the original value proposition of cryptocurrenices such as Bitcoin to link real-world identities to on-chain activities as it would be the recreation of a trusted network, only with higher operational costs.
[34] During the summer, the proof-of-concept Nasdaq Private Market/Exact Equity was essentially a centralized database model and in fairness this is probably more of an exploration of whether this is really a better way to build the mousetrap rather than the final model. See: Nasdaq Launches Enterprise-Wide Blockchain Technology Initiative from Nasdaq.
[35] Nasdaq and Chain to Partner on Blockchain Technology Initiative from Nasdaq.
[36] Nasdaq Unveils Blockchain-Enabled Platform Ling, Announces 6 Inaugural Clients from *Forbes.* Also, personal correspondence.
[37] A new project called T0 announced that it is doing something similar with the Bitcoin blockchain. See Judd Bagley's tweet and story Hedge Fund Borrows $10M in Stock Via the Bitcoin Blockchain from *Wired*. Another startup called Symbiont appears to have done a similar proof-of-concept: Wall Street, Meet Block 368396, the Future of Finance from *Bloomberg.* See also: Nasdaq Blockchain Chief: Bitcoin's Currency and Ledger are 'Two Innovations' by *CoinDesk.*
[38] A blockchain-based property ownership recording system by Alex Mizrahi.
[39] One reviewer noted that: "The regulatory requirements incumbent upon any securities exchange—whether via blockchain or not—make the problems Proof-of-work is intended to solve entirely irrelevant (which is why something like Paxos or PBFT makes more sense, if one decides to use a blockchain at all). On one hand, all parties on-exchange must be "already trusted" in a range of specifically verifiable ways (licenses, charters, accounts, etc.), but something like PoW couldn't prevent things like LIBOR fixing, or really even CDS problems. All of this can be done off-chain and doesn't require network cooption."
[40] While the focus of this paper relates to Bitcoin-like pseudonymous networks, one reviewer explained a potential exception related to Ripple: "In the case of Ripple, market maker accounts will be funded at gateways by traditional settlement ways. They take the form as 'virtual nostro accounts'. Thus, the traditional way supports the new way, specifically Earthport on Ripple (centralized on decentralized). In this case, a pseudonymous network uses a different type of consensus method that is not based on proof of work."
[41] No, Bitcoin is not the future of securities settlement by Robert Sams (2015c).
[42] A few known Bitcoin mining farms by Tim Swanson
[43] Zerocash: Decentralized Anonymous Payments from Bitcoin by Ben-Sasson et al., p. 3; The Bitcoin Backbone Protocol: Analysis and Applications by Garay et al., p. 2.
[44] See Appendix B as well as Some Crypto Quibbles with Threadneedle Street by Robert Sams.
[45] Peer-to-Peer : Harnessing the Power of Disruptive Technologies edited by Andy Oram, Chapter 16.
[46] See The inevitable demise of unconfirmed Bitcoin transactions from Bram Cohen. The largest confirmed double-spend was for $10,000 on OKPay in 2013.
[47] The Ethereum white paper classifies Bitcoin and Bitcoin-like systems as a "state transition system."
[48] The Bitcoin Backbone Protocol: Analysis and Applications by Garay et al., p. 7.
[49] Ibid, p. 9
[50] Spacecoin: A Cryptocurrency Based on Proofs of Space by Park, et al. Thanks to Bram Cohen for pointing to this paper and to Richard Brown for the bond analogy.
[51] See Block Size: Bitcoin Does Not Scale Effectively by Venzen Khaosan and Cost? Trust? Something else? What's the killer-app for Block Chain Technology? by Richard Brown.
[52] In today's terms, to brute force the network — to attack it head on through its hypothetical 'Maginot Line' – would, in theory, cost half of $45,000 per hour (half of the aggregate of 6 blocks as Satoshi suggested above) to obtain the magical "51%" of the hashrate needed to continuously double-spend. In reality, the actual cost is significantly less due to out-of-band, rubber hose attacks and the potential for "selfish-mining." This includes

blatant bribery and hacking of pool. For instance in 2014 a Canadian router was hacked via Border Gateway Protocol, fooling miners. $84,000 in cryptocurrency was stolen. See Hacker makes $84k hijacking Bitcoin mining pool from *The Guardian*. See also: Majority is not Enough: Bitcoin Mining is Vulnerable by Eyal & Sirer.

[53] The myth of a cheaper Bitcoin network: a note about transaction processing, currency conversion and Bitcoinland by Tim Swanson.

[54] Tezos: A Self-Amending Crypto-Ledger Position Paper by L.M. Goodman, p. 7.

[55] From Bram Cohen, starting at 48m45s: Removing the Waste from Cryptocurrencies: Challenges and More Challenges

[56] This energy consumption is on par with a large city, or in one study, the size of a small country: Ireland. See The flow of funds on the Bitcoin network in 2015 and Bitcoin Mining and its Energy Footprint by O'Dwyer and Malone. It also bears mentioning that over time, as the money supply tapers off (via its scheduled decay formula), that the amount of energy used to produce proofs-of-work could also decline but that is not certain.

[57] See Appendix B in Consensus-as-a-service by Tim Swanson.

[58] Public versus Private Blockchains Part 1, Permissioned Blockchains and Public versus Private Blockchains Part 2, Permissionless Blockchains from Bitfury. See also: A dissection of two Bitfury papers by Tim Swanson

[59] It also bears mentioning that this example is for illustrative purposes only, and is not indicative of what those two companies may do in the future. In fact, there are also other startups that have built out a similar multisig and intermediary framework that does not capitalize or utilize Bitcoin.

[60] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder, p. 35.

[61] Founded in 2009, SharesPost was one of the first companies to create a platform for investors to trade private securities investments. It has a partnership with Nasdaq. Similarly, SecondMarket provides a similar venue for trading private company stock and was acquired by Nasdaq in October 2015.

[62] See About eShares and Techcrunch eShares profile. See also EShares, Now Valued At $77 Million, Looks Far Beyond Silicon Valley from *TechCrunch.*

[63] See A cloud-based cap table for startups from *ZDNet.*

[64] Thanks to Robert Sams for this insight. Gemini is a new cryptocurrency exchange that will enable users to "short" a trade. See Bitcoin Exchange Gemini Approved for Launch in New York and The Winklevoss Brothers on Gemini, the 'NASDAQ of Bitcoin'.

[65] We see this "undermining" happen with existing virtual assets: in the fall of 2013 and early 2014, Chinese traders hacked into reporters' Weibo accounts and uploaded fake government documents to spook the market. These hackers sold bitcoins beforehand and after the Weibo account was restored, and repurchased bitcoins at a lower level. Subsequently, we saw continuous effort from traders and cryptocurrency boosters to create sock puppets and use social media (Twitter, Reddit) to spread similar rumors to move the market up and down. These incidents were not reported in Western press, but were relayed by a variety of investors and traders in early 2014.

[66] Bitcoin Blockchain for Distributed Clearing: A Critical Assessment by Robert Sams

[67] Tendermint for Fast Settlement by Jae Kwon.

[68] Satoshi Nakamoto foresaw this noting in the original FAQ, stating that "When Bitcoins start having real exchange value, the competition for coin creation will drive the price of electricity needed for generating a coin close to the value of the coin." There are at least five exceptions to this, described on p.46 in *The Anatomy of a Money-like Informational Commodity* by Tim Swanson.

[69] The fat tail risk is an additional risk when representing an off-chain asset on a chain that requires a token to operate. Systems such as Ripple likely have a similar problem with dependency on XRP. When such a discussion is brought up, the risk is dismissed as being very unlikely to happen, which is the definition of a fat tail event: a very unlikely event that can still have a significant impact. See also: Tokens and Tails by Ayoub Naciri and Elementary Statistics of Tail Events by William Nordhaus,

[70] Based on these assumptions and set of incentives, the goal is to create what (Luu et al. 2015) describe as a "consensus computer." See Demystifying incentives in the consensus computer by Luu et al., p.1; Tendermint: Consensus without Mining by Jae Kwon, p.2; and The Economics of Bitcoin Mining or, Bitcoin in the Presence of Adversaries by Kroll et al.

[71] As described on p. 5 (Park et al. 2015): "The miners in a cryptocurrency are strategic agents who seek to maximize the reward that they get for mining blocks. As such, it is a crucial property of a cryptocurrency that "following the rules" is an equilibrium strategy: in other words, it is important that the protocol rules are designed

in such a way that miners never find themselves in a situation where “cheating" and deviating from the rules yields more expected profit than mining honestly.”

[72] Bitcoin Mining Calculator and Profitability Calculator from CoinWarz

[73] In practice, laborers on the Bitcoin network must account for the capital costs of their hashing equipment, rent for the land, administrative overhead, taxes and – increasingly importantly – the energy costs which can be very specific to their locality, depending on the equipment’s geographic location. All of these costs are tallied against an inelastic wage which can only be attained if the hashing equipment they control is able to outcompete other such miners. Each block race is a zero-sum game, as there is no simultaneous prize for second place. Hashrate is not the exclusive factor in “winning;” network connectivity is also important. In practice – and in the aggregate – depending on its hashrate, the next largest mining pool can potentially obtain a prize (block reward). This is not shared with other pools.

[74] As it lacks any fundamental value, there is no known model that accurately predicts long-term prices. See What is the “real” price of bitcoin?

[75] The math that yields that number is as follows: $300 x 25 bitcoins x 6 blocks per hour equals $45,000 per hour. The caveat: pre-sold contracts where the cloud hashing entity is locked in to mining, even when the bitcoins mined are less valuable then the electricity cost. The most common way out of this is the bankruptcy of the entity. Another tangential caveat: where voting influence is valued, some actors looking to remain anonymous, may be prepared to pay a higher premium.

[76] One reviewer observes that “the watermarking approach or whatever color coins is… suffers from Gresham's syndrome (as in Gresham's law - two units tied together artificially). In the colored coin sense, any colored coin has both the asset claimed value and the bitcoin value. In this sense, you can see the fee rising and the various price scenarios having an impact. If BTC goes to $1,500 that means the 10,000 satoshi comprising a colored coin might impact things. Not only the pricing effects but also the customary approach taken by miners to analyze the transaction based on fee size and make decisions using metrics and methods oriented to pure BTC transactions.”

[77] Overview of Colored Coins by Meni Rosenfeld, p. 2.

[78] A blockchain-based property ownership recording system by Alex Mizrahi.

[79] The nominal fee structure for Omni can viewed on p.4 of its white paper: Omni Protocol Specification (formerly Mastercoin). For a discussion on seigniorage as it relates to Bitcoin see: The Marginal Cost of Cryptocurrency by Robert Sams

[80] Future analysis should do a cost comparison for existing real-world asset tracking systems. For instance, what would be the economic cost to attack DTCC How much would it cost to bribe all the system administrators and security guards?

[81] CoinGecko is currently tracking 244 virtual cryptocurrencies. See also Ray Dillinger’s “necronomicon” – list of dead alt coins.

[82] “Because every transaction published into the blockchain imposes on the network the cost of needing to download and verify, there is a need for some regulatory mechanism, typically involving transaction fees, to prevent abuse. The default approach, used in Bitcoin, is to have purely voluntary fees, relying on miners to act as the gatekeepers and set dynamic minimums. This approach has been received favorably in the Bitcoin community, particularly because it is "market-based", allowing supply and demand between miners and transaction senders to determine the price. The problem with this line of reasoning is, however, that transaction processing is not a market; although it is intuitively attractive to construe transaction processing as a service that the miner is offering to the sender, every transaction that a miner includes will need to be processed by every node in the network, so the vast majority of the cost of transaction processing is borne by third parties, and not the miner that is making the decision of whether or not to include it.” Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform, p. 20. Empirically, this also happens on other proof-of-work chains, such as Dogecoin via Dogeparty.

[83] Overview of Colored Coins by Meni Rosenfeld, p. 5.

[84] August 23rd 2015 Network Statistics from Organ of Corti.

[85] See July 2015 flood attack, Part 5 – Stress Test Analysis from *TradeBlock* and Bitcoin Network Still Backlogged With Tens of Thousands of Unconfirmed Transactions, Causing Delays from *Bitcoin Magazine.*

[86] Slicing data: what comprises blockchain transactions? By Tim Swanson.

[87] Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform, p. 10.

[88] Demystifying incentives in the consensus computer by Luu et al.

[89] Source: Bitnodes.
[90] The drop over the past year was actually around 12% until early October 2015, when a series of malleability and spam attacks took place, creating very large mempools which many nodes were unable to support, forcing them offline. See Bitcoin Node Numbers Fall After Spam Transaction 'Attack' from *Coindesk*, Ongoing Bitcoin Malleability Attack (Low S / High S) from CoinKite, The Who, What, Why and How of the Ongoing Transaction Malleability Attack from *Bitcoin Magazine*, and I Broke Bitcoin from *Motherboard.*
[91] In March 2014, there were approximately 10,000 reachable nodes. See What Are Bitcoin Nodes and Why Do We Need Them? from *CoinDesk.*
[92] According to calculations from Rusty Russell, during this timeframe the average download and upload speeds in both the US and UK measurably increased, so it is unclear what specific factor(s) have led to this decline. It was likely a combination of factors including bandwidth, disk space, and hard caps on residential bandwidth. See Broadband Speeds, New Data.
[93] Dave Hudson has raised the point several times that while hashrate has increased, block making itself remains relatively centralized. See Insights From The Data Mine And Other Adventures Around The Block from *Epicenter Bitcoin* and Dave Hudson explains Bitcoin mining hash rate statistics. See also Is Bitcoin a Decentralized Currency? by Gervais et al.
[94] Integrating, Mining and Attacking: Analyzing the Colored Coin "Game" by Ernie Teo.
[95] Slicing data: what comprises blockchain transactions? by Tim Swanson.
[96] In mid-August 2015, node count dramatically increased by 50% – from 6,000 to 9,000 – however, it was quickly discovered that these were fake nodes.
[97] See A Fast and Scalable Payment Network with Bitcoin Duplex Micropayment Channels by Decker and Wattenhofer, The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments by Joseph Poon and Thaddeus Dryja, and Let's Talk Bitcoin! #242 - Ride The Lightning Network with Rusty Russell.
[98] One explanation for why Luke-Jr. blocks transactions related to SatoshiDice: "SatoshiDice's misuse is *messaging*, not real commerce."
[99] A relatively small transaction fee would not necessarily incentivize miners to include watermarked tokens in the long run. As described in (Teo 2015), miners may attempt to hold-up the network and charge significantly higher fees due to the new exogenous value that is being secured.
[100] This example originally comes from Can Bitcoin's internal economy securely grow relative to its outputs? by Tim Swanson. I would like to thank Antony Lewis and Arthur Breitman for their thoughts and comments on this section.
[101] As noted by one reviewer: "Payoff from unwinding trades is only one possible payoff. Surely, the market for something settled mostly via the Bitcoin blockchain will itself suffer from a chain reorg, creates uncertainty in all the assets the chain clears. Would expect value of the assets backing the colored coins to go down and volatility to go up. That's another payoff vector: 1. short the color market (or buy color volatility), 2. short BTC to hedge the coinbase reward you'll earn during the attack, and 3. buy a ton of hashing power and attempt a reorg."
[102] One reviewer explained that: "With most financial transactions the implications are usually related to the deltas between two relative outcomes. The question really seems to be one of the timing of trades. Consider that Apple Inc. is worth, say, $670B. Now let us say that they issue a new quarterly earnings statement that says that the iPhone 6S blows up after 6 months and they have to do a recall that will cost them $20B. Their share price tanks by 10%. They just lost $67B. Now consider that I can attack the transaction set; I can let some through and block others by re-orging and re-mining blocks. As a large stock holder who owns, say 1% of Apple's stock my stock holding is going to very quickly be worth $670M less than it was; how much is it worth for me to a) block everyone else's trades, b) ensure that mine get serviced and not re-orged away. Say I did get my trade through, what's it now worth 2 days later when the price has dropped 20% to not have a re-org that unwinds my trades? The same problem also exists on upside news. Now if we throw derivatives trading in here then things really can become a mess. That 10% shift in stock price can have much more dramatic consequences then."
[103] As one reviewer pointed out: "The incentive to attack is a simple expected value calculation on the probability of success. Modelling that probability is where it gets complicated. However, you hardly need to embark on that exercise as some simple back-of-envelope calculations make it clear that proof-of-work is a poor security model if there is a meaningful payoff to a double-spend (or even the second-order consequences of a double-spend); something that pseudo-anonymous bearer assets via colored coins provide if they ever existed in meaningfully large value. Even a credible market for shorting BTC gives you the right payoff vector. But what about the BTC

market itself declining with this news, won't that incentivize miners not to attack? This argument assumes that there is no way to hedge the long BTC exposure of the attacker. If there is a credible marketplace for shorting, the attack strategy can be modelled elegantly, no assumptions need be made about the impact of the attack on BTC price whatsoever. It is a trade, no different in principle from hundreds of other market manipulation strategies like squeezing a bond auction or something. Only difference is: Bitcoin does not have a mechanism for regulating market practice. Of course, regulators could prosecute an attacker (if he could be identified) under one of the market manipulation regulations. But then we arrive at the supreme irony that proof-of-work "security" ultimately rests on regulating the cryptocurrency market itself." See also: Creative angles of attacking proof-of-work blockchains

[104] It bears mentioning that at the time of this writing there has not been a publicly known situation where a rogue miner deliberately forked blocks out of the chain for the purposes of double spending.

[105] Script is the built-in scripting language used for creating and customizing transactions. See Demystifying incentives in the consensus computer by Luu et al., p. 1.

[106] In (Luu et al. 2015) the authors were looking specifically at Ethereum as it relates to this issue. However, they noted that there is a verifier's dilemma in Bitcoin: "As the number of transactions increases per block on the Bitcoin network, the 'common good' work required to verify all the transactions in a block approaches a nontrivial quantity. Indeed, some Bitcoin miners have already decided not to verify transactions. On July 4 this year, a serious incident on the Bitcoin network was reported, wherein large pools extended blocks containing invalid transactions [25]. These pools mined on blocks without first verifying the block's transactions, causing a fork in the blockchain." Demystifying incentives in the consensus computer by Luu et al., p. 5.

[107] Hidden surprises in the Bitcoin blockchain and how they are stored: Nelson Mandela, Wikileaks, photos, and Python software by Ken Shirriff.

[108] Bitcoin Hurdles: the Public Goods Costs of Securing a Decentralized Seigniorage Network which Incentivizes Alternatives and Centralization by Tim Swanson, p. 28-30.

[109] Core Development Update #5 by Gavin Andresen.

[110] Make Master Protocol harder to censor #248 by Ron Gross and the Bitcoin Talk debate between Luke-Jr. and Ron Gross.

[111] Update `OP_RETURN` Size Limit to 80 Bytes #690 by Adam Krellenstein and Change the default maximum OP_RETURN size to 80 bytes #5286 by Flavien Charlon.

[112] Starting at 6m24s: Removing the Waste from Cryptocurrencies: Challenges and More Challenges

[113] As one reviewer noted, "If Bob really wanted to obfuscate commitments into a layer 2 system, Bob could pretty much just make a colored coin transaction look just like a standard multisig transaction. For instance Bob could pay to Hash(Pubkey) and Hash(Pubkey+Hash(nonce || data))). Then he would publish nonces into the layer 2 system and thus he could prove his coin has been secretly colored."

[114] Block Size Increase from BitFury Group, p. 2.

[115] Overview of Colored Coins by Meni Rosenfeld, p. 4.

[116] David Schwartz provides an excellent explanation on How does merged mining work?

[117] In terms of charity and altruism, (Garay et al. 2015) believe that: "In this way our results flesh out the incentive compatibility problems of the Bitcoin backbone, but (on a more positive note) they also point to the fact that honest hashing-power majority is sufficient to maintain the public ledger (under favorable network conditions), and hence suggest that the Bitcoin protocol can work as long as the majority of the miners *want it to work* (without taking into account the rationality of their decision)," p. 33.

[118] Tezos: A Self-Amending Crypto-Ledger Position Paper by L.M. Goodman, p. 7.

[119] Public versus Private Blockchains Part 1, Permissioned Blockchains p. 15.

[120] Source: BitInfoCharts of Bitcoin and Namecoin.

[121] This is equivalent to the miners – the labor force – being told they would receive a 50% pay cut. See Entry and Exit Leads to Zero Profit for Bitcoin Miners from the Federal Reserve Bank of New York.

[122] Feathercoin (case study by MaxMiner) and Auroracoin. It bears mentioning that it is beyond the scope of this paper to look at all potential 51% attack scenarios.

[123] In practice, at least one altcoin called Coiledcoin, which attempted to merge mine with Bitcoin, was successfully attacked. CoiledCoin.

[124] Dogecoin Community Celebrates as Merge Mining with Litecoin Begins from *CoinDesk.*

[125] Source: BitInfoCharts of Litecoin and Dogecoin

[126] (Saberhagen 2013) discussed Bitcoin's predetermined emission rate (his term for a staggered increase in the money supply): "The original intention was to create a limited smooth emission with exponential decay, but in fact we a have a piecewise linear emission function whose breakpoints may cause problems to the Bitcoin infrastructure."CryptoNote by Nicolas van Saberhagen, p. 3.
[127] On The Longest Chain Rule and Programmed Self-Destruction of Crypto Currencies by Nicolas Courtois.
[128] On August 25, 2015 the Litecoin network built block 840,000. See Litecoin Network Experiences First Mining Reward Decline from *CoinDesk.*
[129] See Litecoin hashrate analysis post halving by Charlie Lee. One reason that both the Litecoin price and hashrate increased during the summer of 2015 was due to a Ponzi scheme based out of China. See Chinese Promoter Pumping Litecoin Via Ponzi Scheme by Arthur Hayes.
[130] Why did Litecoin's hashrate peak on December 25, 2014? I have reached out and spoken to a number of people familiar with the Litecoin ecosystem including Andrew Vegetabile ("TheRealMage"). His view is that it likely is a trailing effect from the AuxPOW from Dogecoin – that the additional hashrate came from the rolling surge of Dogecoin-based miners.
[131] Why Onename is Migrating to the Bitcoin Blockchain from Onename
[132] Ibid
[133] As explained in (Garay et al. 2015): "We finally note that the Persistence and Liveness properties we put forth and prove should not be interpreted as proofs that all Bitcoin's objectives are met. In particular, they do not guarantee that miners are properly incentivized to carry out the backbone protocol, and they can only offer guarantees in a setting of an *honest majority* amongst a fixed number of players as opposed to a setting where there is an ever changing population of parties acting rationally; see related work below as well as Section 7 for further discussion," p. 5.
[134] "Miners' Incentives and the Decentralized Network" by Ernie Teo.
[135] At the time of this writing, the only public commentary on OFAC from an attorney as it relates to Bitcoin is: What is OFAC and how does it apply to Bitcoin? by Joshua Garcia
[136] The Gambler's Guide To Bitcoin Mining by Dave Hudson and Is Bitcoin a Decentralized Currency? by Gervais et al.
[137] A full list of mining pools and dates can be seen here: What has been the reaction to permissioned distributed ledgers?
[138] For more on the history of BitFury, see Bitcoin miner steps out into the light from *bne IntelliNews,* as well as The 21 companies that control bitcoin from *Business Insider.*
[139] Look inside the surreal world of an Icelandic bitcoin mine, where they literally make digital money from *Business Insider.*
[140] Projects like Confidential Transactions can make the ability to trace, track, or detect the lineage or provenance of a token much more difficult.
[141] Bitcoin-NG stands for Next Generation. See Researchers Tackle Tomorrow's Blockchain Problems With Bitcoin-NG from *CoinDesk.*
[142] I would like Dustin Byington for deriving that term.
[143] KYB stands for Know-Your-Business and KYC stands for Know-Your-Customer. The topic of a "legal entity identifier" (LEI) is beyond the scope of this report.
[144] In theory, Ripple would meet the criteria for "permissionless." In practice, however, the classification is somewhat nebulous and is beyond the scope of this paper.
[145] For Coinbase see Appendix 1: Prohibited Businesses and Prohibited Use and Circle's user agreement.
[146] Bitcoin Network Shaken by Blockchain Fork by Vitalik Buterin.
[147] Another smaller, yet equally expensive fork, took place in early July 2015 – 9 invalid blocks worth around $50,000 of economic activity were mined by a couple pools. See: Some Miners Generating Invalid Blocks from Bitcoin.org and Double Spending Risk Remains After July 4th Bitcoin Fork from *CoinDesk.*
[148] What is Permissioned-on-Permissionless by Tim Swanson.
[149] In their current terms of service, 21inc appears to collect identification information on their users: 21 Terms Of Service, Acceptable Use Policy, and Privacy Policy.
[150] Enabling Blockchain Innovations with Pegged Sidechains p.3 from Back *et al.* For more on Byzantine participants see BAR Fault Tolerance for Cooperative Services by Aiyer *et al.*

[151] Another potential term for DMMS is what (Grigg 2015) called a Nakamoto signature. The Nakamoto Signature by Ian Grigg.
[152] Secure Multiparty Computations on Bitcoin by M. Andrychowiz, S. Dziembowski, D. Malinowski, and L. Mazurek; and How to use bitcoin to design fair protocol by I. Bentov and R. Kumaresan.
[153] Coinbase Seeks 'Invasive' Details on US Bitcoin Mining Operations from CoinDesk.
[154] I spoke to a company representative on August 13, 2015 and they confirmed the statistic first posted on *Business Insider*.
[155] Bitcoin Exchange Highlights to Investors Currency's Ability to Evade Sanctions from *The Washington Free Beacon.*
[156] Coinbase pitch deck from *The Washington Free Beacon.*
[157] Visa, MasterCard stop supporting bank cards in Crimea from *Reuters.*
[158] Bitcoin: A Peer-to-Peer Electronic Cash System by Satoshi Nakamoto.
[159] Bitcoin Exchange Highlights to Investors Currency's Ability to Evade Sanctions from *The Washington Free Beacon.*
[160] A month later, in March 2015, the chief compliance officer resigned from Coinbase, citing personal reasons. See Coinbase Exec Resigns as Company Faces Criticism from *The Washington Free Beacon.*
[161] See Treas., Resource Center, Frequently Asked Questions And Answers.
[162] *The Law of Bitcoin* edited by Stuart Hoegner, p. 201.
[163] This is similar to the issue that Peter Todd highlighted in his report on Ripple: "A closely related issue is what incentive does a node operator have to publicly perform validation services? Ripple currently has no compensation mechanism for public validators, yet validation at minimum raises potential legal issues such as lawsuits for negligence or aiding financial crimes." See p. 6 from Ripple Protocol Consensus Algorithm Review.
[164] On August 27, 2015, a developer on the Bitcoin development mailing list announced a proposal (a BIP) for implementing AML-KYC within Bitcoin, which would enable miners to comply with OFAC regulations. BIPS proposal for implementing AML-KYC in bitcoin by Prabhat Kumar Singh. Singh previously created an altcoin called Voncoin.
[165] Public versus Private Blockchains Part 1, Permissioned Blockchains and Public versus Private Blockchains Part 2, Permissionless Blockchains from Bitfury. See also: A dissection of two Bitfury papers and A few known Bitcoin mining farms by Tim Swanson
[166] There is still the added problem that a trusted or accredited miner still wins a block on a probabilistic basis (an inhomogeneous Poisson process). That means that there is no guarantee for in any given time period that any (or all) of a financial institutions transactions will be processed by the miner they outsource proof-of-work to.
[167] See also p. 25, New York City Bar coursebook: "If I Were a Virtually Rich Man: How the Bitcoin Revolution Will Affect the Banking and Commercial Payments World."
[168] Or in other words, ownership of assets is defined by law, and not just possession of a private key.
[169] One hypothetical involving mining (or block making) in terms of facilitating: if Bob's privkey is stolen and the UTXO(s) associated with the address are spent, there are two possibilities: (1) the UTXO is spent to an address with another UTXO associated with said address; or (2) the UTXO is spent to an address with no other UTXOs associated with it. If miners are informed of the theft beforehand, are they liable for processing these identifiable UTXOs?
[170] The terms "intermediaries" and "facilitators" are not terms of art. For the purposes of this paper, these terms derive their meaning from Technological Change, Financial Innovation, and Financial Regulation in the U.S.: The Challenges for Public Policy by Lawrence White, p. 2-3. See also: Technological Change, Financial Innovation, and Diffusion in Banking by Frame and White.
[171] Mining pools process transactions and get paid by the system: they are effectively payment service providers. An "agent" holds something (like a title) in trust, however having title means you have ownership and control. Miners do not have this control. They do not have the ability to breach the trust relationship by controlling the actual title (e.g., UTXOs) although they have the ability to censor, block and double-spend transactions.
[172] FIN-2014-R008 from the Financial Crimes Enforcement Network.
[173] EBA Opinion on 'virtual currencies', p. 5.
[174] The Bitcoin Backbone Protocol: Analysis and Applications by Garay et al., p. 2.
[175] Regarding settlement, as (Schroeder 2015) points out: "More importantly, "settlement" does not merely include transfer of a security from the seller to the buyer. It also includes the transfer of the purchase price from the buyer to the seller. Consequently, even if the cryptosecurity itself were to be transferred on the blockchain,

settlement could not be completed immediately unless payment was also made over the blockchain." See p. 50 Bitcoin and the Uniform Commercial Code by Jeanne Schroeder.

[176] With respect to payments, according to FinCEN, irrevocability may not be a needed feature in all financial IT systems. For example, "Liberty Reserve also is a completely irrevocable payment system and digital currency. The fact that the transactions are irrevocable, meaning that they cannot be reversed or refunded in the event of fraud, makes it a highly desirable system for criminal use and a highly problematic one for any legitimate payment functions. Revocability protects merchants and users from fraud and is a common feature of legitimate payment systems." See Notice of Finding That Liberty Reserve S.A. Is a Financial Institution of Primary Money Laundering Concern from FinCEN

[177] Bitcoin Blockchain for Distributed Clearing: A Critical Assessment by Robert Sams

[178] For instance, if, as many venture capitalists have openly requested, governments and banks buy hashing equipment and begin mining on the Bitcoin network, only a matter of capital and thermodynamics (exergy) is required to reverse transactions. One definition of exergy is, "In thermodynamics, the exergy of a system is the maximum useful work possible during a process that brings the system into equilibrium with a heat reservoir." See also Bitcoins: Made in China. In June 2015, Fred Wilson stated that banks should run mining equipment: Banks and Brokerages Should Be Mining The Blockchain; in early August, William Mougayar emphasized that banks should buy mining equipment: Dear Big Bank CEO, Re: Blockchains: Obliterate, don't Automate; in early March, Vinny Lingham predicted that the biggest "buyers of bitcoin" in the next 3-5 years will be government: The inaugural Bitcoin Africa Conference 2015 in Cape Town, South Africa.

[179] Banks and Brokerages Should Be Mining The Blockchain by Fred Wilson.

[180] As explored in Appendix B, if a bank intended to become a colored coin issuer as a means of preventing double-spend risks, the operating costs would multiply by several orders in magnitude. There is no free lunch in mining.

[181] Public versus Private Blockchains Part 1, Permissioned Blockchains from Bitfury Group, p. 15

[182] (Narayan 2015) implies that censorship resistance is the ultimate utility for any kind of *real* blockchain, and every other needed feature can be obtained by using a replicated database. This is a superficial understanding of the problem set facing institutions which serve capital markets, and it seems to devolve into a No True Scotsman fallacy. As Simon Taylor has pointed out, permissioned blockchains whose validators are legally known, contract-bound entities solve a separate set of use-cases than unpermissioned blockchains currently can. See "Private blockchain" is just a confusing name for a shared database by Arvind Narayanan.

[183] EBA Opinion on 'virtual currencies', p. 39.

[184] The concept of governance authority is derived from the European Central Bank, *Harmonised oversight approach and oversight standards for payment instruments*, February 2009. There, the governance authority is described as being accountable for the overall functioning of the scheme that promotes the (initiation of the) payment instrument in question, and for ensuring that all the actors involved comply with the scheme's rules. Moreover, it is responsible for ensuring the scheme's compliance with oversight standards.

[185] Ibid

[186] TrustGuard: Countering Vulnerabilities in Reputation. Management for Decentralized Overlay Networks by Mudhakar Srivatsa, Li Xiong and Ling Liu.

[187] A brief summary of their defensive mechanism proposal suggests that: "In TrustGuard, it is assumed that the base overlay network is resilient to attack and provides a means for authenticating messages. Under these assumptions, TrustGuard uses control theory as the basis for its strategic oscillation guard. Using different empirically determined weighting for the guard, the system can mitigate many of the malicious attacks. The fake transaction detection mechanism uses irrefutable proofs to prevent input resulting from fake transactions from being admitted into the reputation system. Assuming each entity in the network has an associated public/private key pair, a transaction proof is exchanged for each transaction, allowing claims of malicious activity to be checked by a trusted third party." See: A Survey of Attack and Defense Techniques for Reputation Systems by Kevin Hoffman, David Zage and Cristina Nita-Rotaru, p. 26.

[188] An Efficient Distributed Currency by Ben Laurie.

[189] See Fedcoin by JP Koning, Fedcoin: On the Desirability of a Government Cryptocurrency by David Andolfatto, A Central Bank "cryptocurrency"? An interesting idea, but maybe not for the reason we think by Richard Brown, Which Fedcoin? by Robert Sams and Fedcoin—how banks can survive blockchains by Robin Winkler.

[190] Centrally Banked Cryptocurrencies by George Danezis and Sarah Meiklejohn.

[191] Blockchain Finance by Robert Sams and Consensus-as-a-service by Tim Swanson.

[192] Negotiability, Property, and Identity by James S. Rogers.
[193] Regarding book-entry systems see also p. 11 of "Beyond the hype: Blockchains in capital markets" from McKinsey & Co. (forthcoming).
[194] See Coogan, Article 9-An Agenda for the Next Decade, 87 Yale L.J. 1012, 1037-38 (1978).
[195] Negotiability, Property, and Identity by James S. Rogers, p. 480.
[196] Overview of Colored Coins by Meni Rosenfeld, p. 6.
[197] As discussed by Robert Sams: "Blockchain 2.0 is instead an ideological thesis devoted to bringing back bearer securities in cryptographic form. This may be desirable as a political programme for making property titles resistent to corporate and state censorship. But a network that doesn't provide a one-to-one correspondence between what the network and the law say is who-owns-what is a network that can't exist without the very legacy settlement framework that it seeks to replace, for the latter will remain the authoritative record of ownership. Blockchain 2.0 is useless as a solution for financial settlements in cash, securities, and other off-chain property titles." Private mailing list, January 23, 2015.
[198] Both the Bitcoin whitepaper and code announcement were released on a cypherpunk mailing list; and many of the initial conversations revolved around ways to route around governments and financial institutions.
[199] An argument could be made that with ZKSNARKs and "confidential transactions" a type of "censorship-resistance" could be enabled, however that is topic for a different paper. There are specific problems that different types of blockchains are attempting to solve. However, the utility and desirability of non-public blockchains such as quasi-private consortium chains is not within the scope of this paper. Readers interested in this topic should survey recent writings on this topic including (Breitman 2015a, 2015b; Brown 2014, 2015b; Buterin 2015; and Greenspan 2015). See What do blockchains accomplish? and A functional nomenclature of cryptographic ledgers by Arthur Breitman; On Public and Private Blockchains by Vitalik Buterin; Ending the bitcoin vs blockchain debate by Gideon Greenspan; A simple model to make sense of the proliferation of distributed ledger, smart contract and cryptocurrency projects and How to explain the value of replicated, shared ledgers from first principles by Richard Brown.
[200] Consumers' Research compilation.
[201] Theft of currency is an issue that has been around since the birth of cash. It is not new to Bitcoin. The difference here however is that Bitcoin is not legal tender and can become encumbered. See also *Nemo dat quod non habet*
[202] BitPay Sues Insurer After Losing $1.8 Million in Phishing Attack from *CoinDesk.*
[203] How the Feds Can Take Even Legally Earned Bitcoins from *Bloomberg*
[204] Bitcoin's lien problem from *Financial Times* and Uniform Commercial Code and Bitcoin with Miles Cowan.
[205] See: Learning from the past to build an improved future of fintech and Too Many Bitcoins: Making Sense of Exaggerated Inventory Claims. Discussion surrounding what this type of relationship are can be found on p. 24, 31 and 44 in Bitcoin and the Uniform Commercial Code by Jeanne Schroeder and p. 178-193 in *The Law of Bitcoin* edited by Stuart Hoegner.
[206] Do the holders of a watermarked token actually hold title, or is it merely an accounting tool used to reflect ownership in the real world? And how might this vary from jurisdiction to jurisdiction? Regarding settlement finality and title transfer see also: Proposal of a DIRECTIVE OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL amending the Settlement Finality Directive and the Financial Collateral Directive by the Commission of the European Communities and T2S Functional Design from Functional Coordination Group.
[207] One reviewer hypothesized that: "It would be interesting to consider this in the context of the "title by registration" concept that Torrens Title land has. The register (or in this case blockchain) IS title. This is a statutory legal construct and so is difficult to apply globally, but if a government wanted to create a legal framework for financial instruments traded on a distributed ledger, it would be a relevant legal concept to explore."
[208] Starting at 32m40s: Bitcoin Regulation panel during the July 2015, *American Banker* regulation panel, this question was asked.
[209] Regarding the finality of settlement with respect to digitization of assets, see p. 12 of "Beyond the hype: Blockchains in capital markets" from McKinsey & Co. (forthcoming)
[210] As one reviewer noted: "The only thing that discharges a monetary obligation absolutely as opposed to conditionally is legal tender. This is the nature of legal tender: if you owe me $3,000 and you deliver $3,000 in cash, that discharges your monetary obligation to me. If you deliver $3,000 in bitcoin, it may or may not. I can still reject the $3,000 in bitcoin you give to me because it is not legal tender. It does not necessarily, by law, discharge

a monetary obligation. If I want to sell my house in exchange for a 1965 Mustang, if you deliver a Mustang with a lien on it, I don't have to accept it. Settlement finality is really about whether the obligation has been discharged or not."

[211] The UCC and Bitcoins: Solution to Existing Fatal Flaw by George Fogg.

[212] According to UCC § 9-102. DEFINITIONS AND INDEX OF DEFINITIONS: (42) "General intangible" means any personal property, including things in action, other than accounts, chattel paper, commercial tort claims, deposit accounts, documents, goods, instruments, investment property, letter-of-credit rights, letters of credit, money, and oil, gas, or other minerals before extraction. The term includes payment intangibles and software.

[213] UCC §9-315(a).

[214] The UCC and Bitcoins: Solution to Existing Fatal Flaw by George Fogg.

[215] Bitcoin's lien problem from *Financial Times.*

[216] The UCC and Bitcoins – Solution to Existing Fatal Flaw by George Fogg.

[217] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder.

[218] Would the securities intermediaries be sure they are shielded from liability to an earlier secured party or adverse claimant? Would a law firm give an opinion to that effect? At the first step in this analysis, the securities intermediary is a transferee of a general intangible. 8-116 doesn't seem to help because it only applies when a securities intermediary receives a financial asset, which is generally the case with the deposit of securities (i.e.,Dabbah Secs. Corp. v. Croesus Capital Corp., 297 A.D.2d 531 (N.Y. App. Div. 1st Dep't 2002)). Here it receives a general intangible, then converts it to a financial asset after receipt.

[219] According to Miles Cowan: "I do wonder, however, whether a securities intermediary would agree to participate in this alchemy. Would the securities intermediary be sure they are shielded from liability to an earlier secured party or adverse claimant? Would you give an opinion to that effect? At the first step in this analysis the securities intermediary is a transferee of a general intangible. 8-116 doesn't seem to help because it only applies when a securities intermediary receives a financial asset, which is generally the case with the deposit of securities (i.e.,Dabbah Secs. Corp. v. Croesus Capital Corp., 297 A.D.2d 531 (N.Y. App. Div. 1st Dep't 2002)). Here it receives a general intangible then converts it to a financial asset after receipt." See also: Is UCC Article 9 the Achilles Heel of Bitcoin? and Is UCC Article 8 Bitcoin's Savior (for Commercial Law)? by Bob Lawless.

[220] It could be the case that several years down the line, many of the laws and regulatory frameworks cited in section 3 are amended to take into account certain characteristics of virtual currencies. For instance, it will be worth revisiting if bitcoins are defined as legal tender by a large jurisdiction.

[221] Cryptocurrencies as Money Under the UCC by Miles Cowan.

[222] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder, p. 56.

[223] Thanks to Robert Sams for providing this thought experiment.

[224] FIN - 2015 - R001 from the Financial Crimes Enforcement Network.

[225] FinCEN Rules Commodity-Backed Token Services are Money Transmitters from *CoinDesk.*

[226] Ibid

[227] According to Juan Llanos, providing this type of service may now require a license: "In the last few rulings, FinCEN has slowly been casting a wider net to subsume virtually every business activity under the money transmission category. I see this ruling as a harbinger for the future inclusion in the definition of money transmission of the last activities that so far had claimed an exemption, namely all software wallets (e.g., Blockchain.info) or services providing a user-friendly interface on top of the blockchain. And this, of course, would open the door to this type of companies to be considered money transmitters at the state level, with the attendant obligation to obtain licenses. Of course, this is speculation at this point, but the phrase "allowing unrestricted transfer of value" could easily be interpreted broadly to include any software. In Llanos analysis, the latest FinCEN ruling, sends new signals to the industry: 1) Brokers that only provide delivery, communication, or network data access services (and do not collect fees from transaction participants) are not money transmitters; 2) Anyone that provides a "manifestation" of ownership or value (in any format) and allows the "unrestricted transfer of value" is a money transmitter. Similarly, electronic trading in e-currencies or e-precious metals are manifestations of ownership or control of real currencies, commodities or precious metals via: 1) Digital certificates of ownership (digital certificate = virtual currency); 2) Paper ownership certificates; 3) Account statements or any other form." Personal correspondence, August 16, 2015.

[228] The overall legal and regulatory framework requires continued exploration and research. On September 17, 2015, the U.S. Commodity Futures Trading Commission (CFTC) announced that it was classifying bitcoin and other

virtual currencies as a "commodity" and therefore was under the preview of its jurisdiction. While this directly impacts bitcoin-related derivatives, it is unknown how this impacts colored coin-related projects that rely on using bitcoin as a carrier layer for off-chain financial assets. See Bitcoin Is Officially a Commodity, According to U.S. Regulator from *Bloomberg,* CFTC Orders Bitcoin Options Trading Platform Operator and its CEO to Cease Illegally Offering Bitcoin Options and to Cease Operating a Facility for Trading or Processing of Swaps without Registering from CFTC, Bitcoin as a Commodity: What the CFTC's Ruling Means from *CoinDesk* and The CFTC's Not-So-Hidden Message: Traders Beware from *CoinDesk.*

[229] Proof-of-existence is typically used to demonstrate that a specific electronic document existed at or before a certain point in time. The original proof-of-existence demonstrations weren't designed for the transfer of titles however over the past year it has been discussed at conferences and social media as a possibility.

[230] GLBSE is a now defunct virtual "stock market" that enabled bitcoin users to purchase, trade and acquire "shares" in a variety of listed companies.

[231] Bitcoin and the Uniform Commercial Code by Jeanne Schroeder.

[232] U.S. Agencies to Say Bitcoins Offer Legitimate Benefits from *Bloomberg.*

[233] See SEC vs Trendon Shavers; SEC's settlement with SatoshiDice; Fraud Allegations Haunt Former CEO of HashingSpace from *CoinDesk*, SEC Seizes Assets from Alleged Altcoin Pyramid Scheme from *CoinDesk*, Feds raid cryptocurrency startup accused of scamming $32M from investors from *ArsTechinca*, and SEC Sues Brother of GAW Miners CEO Amid Investigation from *CoinDesk.*

[234] 15 U.S. Code § 77b - Definitions; promotion of efficiency, competition, and capital formation. See section (a)(1).

[235] One reviewer explained that: "While it is sometimes assumed that if something isn't an investment contract it isn't a security, but that is not true. The "investment contract" characteristic is sometimes thrown around like it is "the test" however certain types of CDs are types of securities. Even if they pay 0%. Yet they don't meet the *Howey* test definition and would still be securities."

[236] *The Law of Bitcoin* edited by Stuart Hoegner, p. 203.

[237] Bitcoin Bucket Shop Kicks Bucket from *Bloomberg.*

[238] See Mitigating the Legal Risks of Issuing Securities on a Cryptoledger and 18 Months of Token Presales: A Review by Tim Swanson.

[239] SEC v. Howey Co., 328 U.S. 293 (1946).

[240] Recall that the *Howey* test is not the test for the definition of a security. It is a test for an investment contract. One of the many instruments that are included in the definition of a security. According to one reviewer "if the question is whether a court would find that an altcoin was a security, I am certain I could construct one that meets the definition. I am not certain that I could construct one that I was equally confident failed to meet the definition. In any case, courts do afford rulings by administrative agencies great deference, they do not afford the same to absence of rulings." Note: based upon the *Chevron* deference, a government agency's (such as the SEC) interpretation of a statute can change over time and as of this writing it has not weighed in with the aforementioned scenario ("initial coin offerings" as the equivalent of an unregistered securities offering).

[241] SEC Adopts Rules to Permit Crowdfunding from the Securities and Exchange Commission

[242] Integrating, Mining and Attacking: Analyzing the Colored Coin "Game" from Ernie Teo

[243] See How To Crowd Fund On Multisig - Spells of Genesis Raised Over 860 Bitcoins! from Coinkite.

[244] Personal correspondence, July 12, 2015.

[245] A relevant paper that explores the automation, formalization and digitization of legal (financial) contracts is: Contract as Automaton: The Computational Representation of Financial Agreements by Flood and Goodenough. See also: 'Smart' derivatives can cure XVA headaches by Massimo Morini and Robert Sams

[246] Public versus Private Blockchains Part 1, Permissioned Blockchains and Public versus Private Blockchains Part 2, Permissionless Blockchains from Bitfury.

[247] I would like to thank Robert Sams and Arthur Breitman for articulating this issue.

[248] A blockchain-based property ownership recording system by Alex Mizrahi.

[249] Bitcoin's lien problem from *Financial Times.*

[250] The great pivot? Or just this years froth? by Tim Swanson.

[251] What is Dual Integration? by Casey Kuhlman and The Ricardian Contract by Ian Grigg.

[252] CryptoNote by Nicolas van Saberhagen. For more on Tembusu's "TRUST" model see Combining anonymity and KYC in cryptocurrency protocols by Kravchenko and Kristof.

[253] Zerocoin and Zerocash were originally proposed as add-ons to Bitcoin. Due in part to technical reasons (the proofs would be very large) but also because some elements of the Bitcoin core development team did not want to draw attention from law enforcement agencies, the integration idea was nixed. How could Bitcoin be extended to support Zerocash? From the paper, "The second approach is to extend Bitcoin's scripting language by adding an opcode that invokes VerifyTransaction, with the requisite arguments embeded alongside the opcode script. Such transactions must be exempt from the requirement they reference an input (as they are Zerocash transactions are self-contained), and, like coinbase transactions, be able to create bitcoins ex nihilo (to account for v pub). Moreover, while VerifyTransaction is run at the standard point in the Bitcoin transaction processing flow for evaluating scripts, the coin commitments and spent serial numbers are not actually added to CMList (resp., SNList) until their containing block is accepted (i.e., merely verifying a transaction does not have side effects)." Zerocash: Decentralized Anonymous Payments from Bitcoin by Ben-Sasson et al., p. 28.

[254] Bitcoin is on the verge of a constitutional crisis from *Vox.*

[255] The Pure Theory of Public Expenditure by Paul Samuelson; The Logic of Collective Action by Mancur Olson, Jr. and A Group Is Its Own Worst Enemy by Clay Shirky.

[256] Some of these issues were predicted in a formal risk analysis from The Bitcoin Foundation: Bitcoin Risk Management Study Spring 2014. See two specific sections: "Threats to Decentralization - Software Development" and "Threats to Decentralization – Mining."

[257] See also an interview with Robin Hanson which discusses this governance challenge as it relates to public blockchains: Futarchy, Prediction Markets And The Challenge Of Disruptive Technology from Epicenter Bitcoin

[258] The five with commit access are: Gavin Andresen, Jeff Garzik, Greg Maxwell, Pieter Wuille and Wladimir J. van der Laan. There are several other "core" developers involved in the debate including Mike Hearn, Peter Todd, Adam Back, Matt Corallo, Meni Rosenfeld and Mark Friedenbach.

[259] See the visual-aid: "Historical Chart of Average Blocksize and Network-Imposed Blocksize Limits" as well as [PATCH] increase block size limit from Jeff Garzik, What are the arguments for and against the increase of the block size limit? at Stack Exchange and Sizing up the block size debate by Adamant Research.

[260] Information Propagation in the Bitcoin Network by Decker and Wattenhofer, Creating a decentralised payment network: A study of Bitcoin by Jonathan Levin; Eclipse Attacks on Bitcoin's Peer-to-Peer Network by Heilan et al.; The Bitcoin Backbone Protocol: Analysis and Applications by Garay et al.; and the Ethereum white paper, p.21. Specifically, "There is another factor disincentivizing large block sizes in Bitcoin: blocks that are large will take longer to propagate, and thus have a higher probability of becoming stales. In Ethereum, highly gas-consuming blocks can also take longer to propagate both because they are physically larger and because they take longer to process the transaction state transitions to validate. This delay disincentive is a significant consideration in Bitcoin, but less so in Ethereum because of the GHOST protocol; hence, relying on regulated block limits provides a more stable baseline."

[261] See Demystifying incentives in the consensus computer by Luu et al., p. 5; Adam Back Says the Bitcoin Fork Is a Coup from *IEEE Spectrum*, Majority is not Enough: Bitcoin Mining is Vulnerable by Eyal and Sirer, Game-Theoretic Analysis of DDoS Attacks Against Bitcoin Mining Pools by Johnson *et al.*, Eclipse Attacks on Bitcoin's Peer-to-Peer Network by Heilman *et al.* and Accelerating Bitcoin's Transaction Processing Fast Money Grows on Trees, Not Chains by Yonatan Sompolinsky and Aviv Zohar.

[262] What is the blockchain hard fork "missile crisis?" by Tim Swanson and A summary list of all concerns related to not rising the block size by Jorge Timon.

[263] Block Size Increase from BitFury Group, p. 6.

[264] Brief thoughts on the Bitcoin block size debate by Richard Brown.

[265] See Why is Bitcoin forking? and Blocksize Debate At The Breaking Point both by Mike Hearn, the past five posts from Gavin Andresen, this reddit thread as well as Web-Wallet Providers Divided over Andresen's 20 MB Block Size Increase Proposal and Major Payment Processors in Favor of Block Size Increase; Coinkite and BitPagos Prefer BIP 100 from *CoinTelegraph.*

[266] There is even a disagreement as to what a soft versus a hard fork is.

[267] Seven large, VC-funded Bitcoin startups signed a letter in support of this plan. See Bitcoin Backers Propose Solution to Split That Threatened Disruption from *Bloomberg.*

[268] Brief thoughts on the Bitcoin block size debate by Richard Gendal Brown.

[269] At the time of this writing, according to XTnodes.com, there are 639 Bitcoin XT nodes, which represent about 10% of all nodes. There is a way to spoof and create fake XTnodes ("PseudoNodes") so some of these may be false

positives. See also: ~44% of Bitcoin mining hash power is currently voting to support 8 MB blocks. from Jameson Lopp.

[270] On August 18, 2015 the first block supporting BIP 101 (using the Bitcoin XT implementation) was mined. Block number 370434. First block with a version greater than 3 on Reddit.

[271] There are multiple documented instances of individuals creating dozens, and even hundreds, of fake "pseudo nodes" in an effort to modify the Bitcoin XT node numbers. These nodes pretend to be using the BIP 101 standard but are not. It is a form of a Sybil attack. How to run 3,000 completely legit full nodes aka don't trust the node numbers part 2 on Reddit.

[272] Among many threads see: block-size tradeoffs & hypothetical alternatives by Adam Back.

[273] There are at least three different implementations of Lightning Network that are currently being developed. See HashPlex Unveils Lightning Network Implementation from CoinTelegraph.

[274] Major Mining Pools Make a Stand Against Bitcoin XT Fork, Support for BIP 100 Grows from *Bitcoin Magazine.*

[275] Slush's pool recently underwent a DDoS attack and it is commonly believed that it is due to their support for BIP 101.

[276] Block Size Increase Requirements by Chun Wang; Elastic block cap with rollover penalties and It's time for a break: About the recent mess & temporary new rules response by Meni Rosenfeld; Making Decentralized Economic Policy by Jeff Garzik; Tree-chains preliminary summary by Peter Todd; Seven reasons a "vote" for BIP100 is a bad idea from Bridge21; BIP 100 Support Grows; 21% Attack Worries Remain from *CoinTelegraph*; and A few results from the first intentional stress test on a communal blockchain by Tim Swanson.

[277] Block size votes from Bitcoinity.

[278] For a counter-explanation see On block sizes by Mike Hearn

[279] Bitcoin Network Capacity Analysis – Part 2: Macro Transaction Trends from TradeBlock.

[280] A Fast and Scalable Payment Network with Bitcoin Duplex Micropayment Channels by Decker and Wattenhofer, p. 2

[281] Results of the CoinWallet stress test.

[282] Ibid. See also Bitcoin Network Capacity Analysis – Part 5: Stress Test Analysis from TradeBlock.

[283] According to an anecdote by Jeff Garzik, a Bitcoin Core developer, a large financial institution (Fidelity) built out a product that planned to use the Bitcoin network but was unable to turn it on because it would have maxed out the network capacity very quickly. See Jeff Garzik's "chicken and egg" problem

[284] CoinWallet says Bitcoin stress test in September will create 30-day backlog from *IB Times*

[285] One source alleges that because no organization publicly announced a "test" prior to this event, it was therefore an attack. However, it could likely have been an academic research team investigating the impact a large increase in transactions would have on the network. Furthermore, one of the standing assumptions of Bitcoin is that there would be constant attacks due to the fact that it is a public, untrusted network; thus it is unclear why anyone should be expected to announce a "test" or an "attack." See A Bitcoin Spam Attack Post-Mortem: S(LA)ying Alive by Josh Cincinnati. While not directly related to this test, other researchers are looking at how secure the various hosted API services are. See: blockchain.info / spoofed transactions problem / Aug. 4, 2015. Ironically, some elements of the Bitcoin community have appealed to using the Computer Misuse Act of 1990 in the UK to prevent such "stress tests." However, relying on external governing structures to prevent "attacks" on the network arguably seems inconsistent with the original Bitcoin "ethos" and even security assumptions: that malicious attackers would always be active on the network. See Bitcoin Spam Tests 'Could Violate UK Law' from *CoinDesk.*

[286] A Bitcoin Spam Attack Post-Mortem: S(LA)ying Alive by Josh Cincinnati.

[287] See Slicing data: what comprises blockchain transactions? by Tim Swanson. Contrary to the narrative from "Highlanders" (that there can "only be one chain"), permissioned ledgers will actually help reduce at least one security problem on public ledgers as registered assets will instead be placed on permissioned systems, and not bloat public ledgers.

[288] Eclipse Attacks on Bitcoin's Peer-to-Peer Network by Heilman et al.; Tampering with the Delivery of Blocks and Transactions in Bitcoin by Gervais et al.; Discovering Bitcoin's Public Topology and Influential Nodes by Miller et al.

[289] This was also briefly touched on in Saberhagen 2013, p.12: "Heavy calls in a trust-less p2p network may lead to serious vulnerabilities, because nodes are obliged to check every new block's proof-of-work. If a node spends a considerable amount of time on each hash evaluation, it can be easily DDoSed by a flood of fake objects with arbitrary work data (nonce values)."

[290] One reviewer noted that "it would be possible in theory for a non-watermarked competitor to attack Bitcoin-based competitors to take away their market share. They would have the motive and perhaps even the means to do such an attack, and the whole operation could be done rather anonymously." See: Who stands to benefit from a spam attack on the Bitcoin network? from Piotr Piasecki.

[291] This is not idle speculation. See p. 4 "The argument that there can never be more than 21 million bitcoin because 'if a fork raised the cap, then it wouldn't be Bitcoin anymore' isn't very substantive, for Bitcoin is what the consensus says it is." From Tezos: A Self-Amending Crypto-Ledger Position Paper by L.M. Goodman as well as The 21mm BTC Soft Cap by Ryan Selkis and Bitcoin's Ironic Crisis by Bram Cohen.

[292] The term "folklore" is used in Demystifying incentives in the consensus computer by Luu et al., p. 3.

[293] See CryptoNote by Nicolas van Saberhagen (a pseudonym), p. 10; Tezos: A Self-Amending Crypto-Ledger Position Paper by L.M. Goodman, p. 4.

[294] The likelihood of "mutually assured destruction" (MAD) has not been fully quantified. It has been hypothesized that if neither agrees and a real hard fork takes place, a significant minority could lose some value through some type of attack on the network. See What is the blockchain hard fork "missile crisis?"

[295] It could also be a Mexican standoff between: miners, exchanges and darknet markets (DNMs). DNMs still provide a non-insignificant amount of liquidity to certain exchanges and attract "virgin" coins from miners, therefore the operators could, in theory, have some "say" in the adoption process. See: Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem by Soska and Christin. There is typically a market premium for "virgin" bitcoins– that is to say, "fresh" bitcoins that come directly from mining pools as users can prove the provenance was not affiliated with any illicit activity. One such provider is Blocktrail, a Netherlands-based startup that provides a "virgin-coin"-as-a-service and requires KYC and AML of its customers. Announcing BlockTrail Mint: Fresh Bitcoin Delivered from the Mines. from Blocktrail.

[296] Personal correspondence, June 25, 2015.

[297] Ethereum is attempting to "bomb" the chain to switch to proof-of-stake at a later date; see, Progress on 'the bomb.'

[298] What about the Bitcoin "standard?" Another challenge for watermarked token platforms is that the underlying system it relies upon – such as Bitcoin – is still not a real protocol. Real protocols are based on a formalized standard, such as TLS / SSL. After formalization occurs, there are then differing implementations of that standard, such as OpenSSL, which are developed independent of the standard. In the case of Bitcoin, the Bitcoin Core codebase is the standard and, tautologically, the standard is the Bitcoin Core codebase. Any changes to the consensus-critical codebase directly impacts the entire network because the consensus-critical codebase is the network. According to Dave Hudson, chief architect at Peernova, this creates a fundamental problem when attempting to scale govern-less networks such as Bitcoin; not only do all implementations of Bitcoin need to be bug-for-bug compatible with Bitcoin core (which would be difficult to coordinate, assuming Bitcoin's usage grows and actually decentralizes) but unlike Linux, there is no legitimate "decider" as to how this occurs. See episode #90: Insights From The Data Mine And Other Adventures Around The Block with Dave Hudson. This is a known issue that Emin Gün Sirer, Peter Todd and others have pointed out this past year. It is one of the main motivations behind the development of libbitcoinconsensus,a consensus library. Readers may also be interested in proposals such as found in A Protocol for Interledger Payments by Stefan Thomas and Evan Schwartz and An architecture for the Internet of Money by Meher Roy.

[299] In contrast to permissionless ledgers, permissioned ledgers may have clearer governance due to explicit chain-of-command, terms of service, real-world reputations and contractual obligations. However, they also have a different set of challenges, such as trying to secure identity, handle KYC management, and managing and updating a shared ledger between enterprises. This is beyond the scope of this paper.

[300] ScalingBitcoin Montreal 2015 and Crucial Bitcoin Scalability Topics to be Discussed at New Scientific & Academic Workshops from *Bitcoin Magazine.* For additional perspective, see Mike Hearn's response to both the workshop and open letter from some Core developers.

[301] For a variety of views see: Bitcoin faces a crossroads, needs an effective decision-making process by Arvind Narayanan and Andrew Miller; Bitcoin 'Forked' in Controversial Bid to Resolve Scalability Question by *CoinDesk*; A Coin Needs a Board by Ittay Eyal; Bitcoin Is Having an Identity Crisis from *Bloomberg*; Brief thoughts on the Bitcoin block size debate by Richard Brown; Bitcoin Community in Disarray After New Software Introduced from *Motherboard*; A bitcoin civil war is threatening to tear the digital currency in two from *Business Insider*; Bitcoin

Core Devs in 'Civil War' Insist We're Not Getting The Whole Story from *Motherboard*, Bitcoin could split in debate over currency's future from *BBC*; Bitcoin's Noisy Size Debate Reaches a Hard Fork from *The Wall Street Journal.*

[302] Bitcoin is Undergoing a Governance Crisis, not a Block Size Dilemma by William Mougayar.

[303] One possible reason why Satoshi Nakamoto has chosen to remain anonymous is, in part, so he would not be charged with designing and administering an unlicensed financial instrument and platform that does not gather KYC on its users, for failing to implement an AML and BSA-compliant model to prevent money laundering and other possible legal challenges. See FIN - 2013 - G001 from Financial Crimes Enforcement Network and FinCEN Guidance Validates Bitcoin Industry but Targets Satoshi by Dan Friedberg.

[304] In response to Kumar Singh's proposal for incorporating KYC-AML into Bitcoin, Gavin Andresen suggested that centralizes decision making into a specific organization.

[305] While Bitcoin XT may have relatively centralized decision-making, in its current state the developers cannot "kill" or "withdraw" certain UTXOs from circulation so it may not meet the criteria for being deemed an "administrator." In the Bitcoin XT implementation FAQ, the question "how are decisions made" is answered: "Bitcoin XT is a patch set on top of Bitcoin Core. It is currently maintained by Mike Hearn in collaboration with Gavin Andresen and Tom Harding. If a patch seems in line with the principles of the project it will be considered for inclusion. If the developer is willing to assist with rebasing work, that also helps build the case for inclusion. Decisions are made through agreement between Mike and Gavin, with Mike making the final call if a serious dispute were to arise. In the event that such a dispute cannot be resolved this way and an XT developer were to make a fork, this website would provide a discussion and link to the competing version." See Bitcoin XT Frequently asked questions and An XT FAQ by Mike Hearn.

[306] This is comprised of six tweets.

[307] A Survey of Attack and Defense Techniques for Reputation Systems by Kevin Hoffman, David Zage and Cristina Nita-Rotaru, p. 13.

[308] Maze as discussed in Hoffman et al.: "However, even if source data is authenticated using cryptographic mechanisms, self-promotion attacks are possible if disparate identities or a single physical identity acquiring multiple identities through a Sybil attack [Douceur 2002] collude to promote each other. Systems that do not require participants to provide proof of interactions which result in positive reputations are particularly vulnerable to the attack."

[309] Author Attribution in the Bitcoin Blocksize Debate on Reddit by Andre Haynes. Notable false positives included "PrimeDice" and "BitsByDre," which were tipping accounts.

[310] Ibid

[311] Whole Foods Executive Used Alias from *The New York Times.*

[312] Jonathan Levin originally provided this thought experiment. I would like to thank John Whelan for rephrasing and reframing the question.

[313] I would like to thank both Dave Hudson and Jonathan Levin for bringing this analogy to my attention.

[314] Understanding and Influencing Attackers' Decisions: Implications for Security Investment Strategies by Cremonini and Nizotsev.

[315] One common definition of a zero-day vulnerability is, "an undisclosed and uncorrected computer application vulnerability that could be exploited to adversely affect the computer programs, data, additional computers or a network." I would like to thank Dave Hudson for linking this challenge with Bitcoin.

[316] Overview of Colored Coins by Meni Rosenfeld, p. 6.

[317] All Systems Are Insecure Until Proven Otherwise: An Interview with Bruce Schneier from *Motherboard*. See also Turing Lecture by Adi Shamir, slide 8 ("Absolutely secure systems do not exist; To halve your vulnerability, you have to double your expenditure; Cryptography is typically bypassed, not penetrated") and RSA founders give perspective on cryptography from *Network World.*

[318] As of this writing, neither SSRN nor Arxiv have any papers or prints specifically related to colored coins. In fact, only two papers even mention the idea of "colored coins," yet neither delves into the technical or economic distortions.

[319] The Bitcoin Backbone Protocol: Analysis and Applications by Garay et al., p. 2.

[320] Personal correspondence, August 13, 2015. See also: Bitcoin Meets Strong Consistency by Decker et al.

[321] This example originally comes from Will colored coin extensibility throw a wrench into the automated information security costs of Bitcoin? by Tim Swanson.

[322] If the price of bitcoin does not adequately incentivize the miners, then there will be a difference between value of a bitcoin and the network. Some entity would have to step in to compensate for that difference. Whether collective action is sufficient to provide this compensation is currently unknown, but there are coordination problems inherent in this model that would make it difficult.
[323] The alleged Kissinger quote is, "Who do I call if I want to speak to Europe?"
[324] A Survey of attacks on Reputation Systems by Hoffman, et al.
[325] For a simple overview and explanation of this see: The Hold-up Problem by Kala M. Krishna, A simple economic teaching experiment on the hold-up problem by Balkenborg et al.; and Hold-up Problem by Yeon-Koo Che. See also Bilateral monopoly, specifically p.3 in Recent developments of the bilateral monopoly theory: relevance for the food market analysis by Valeria Sodano.
[326] Apple's Supplier Strategy Aims at Reducing Risks by Stephen DeAngelis
[327] Peter Todd briefly touched on this in his discussion of Ripple, noting a possible attack during a Transaction Flood: "A clever attacker wishing to overtly attack the Ripple network can masquerade their traffic as legit economic activity. From the point of view of Ripple Labs - a major XRP owner - such an attack may not even be considered an attack!" And in a related footnote: "Similar to the debates in the Bitcoin ecosystem about whether or not fee-paying uses of the Bitcoin blockchain for non-Bitcoin-denominated transactions constitute "attacks". See p. 9 Ripple Protocol Consensus Algorithm Review.
[328] OPEC Cheats Most Since 2004 as $100 Oil Heralds Supply from *Bloomberg*, The incentive to cheat: An empirical analysis of OPEC by James Griffin and Split by Infighting, OPEC Keeps a Cap on Oil from *The New York Times.*
[329] Integrating, Mining and Attacking: Analyzing the Colored Coin "Game" by Ernie Teo.
[330] Ibid
[331] Prospect Theory: An Analysis of Decision under Risk by Daniel Kahneman and Amos Tversky.