

Durham Research Online

Deposited in DRO:

30 October 2017

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Down, K. and Hutchinson, M. (2015) 'Bitcoin will bite the dust.', *Cato journal.*, 35 (2). pp. 357-382.

Further information on publisher's website:

<https://www.cato.org/cato-journal/springsummer-2015>

Publisher's copyright statement:

Copyright © Cato Institute. All rights reserved.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

BITCOIN WILL BITE THE DUST

Kevin Dowd and Martin Hutchinson

Bitcoin is the most radical innovation in the monetary space for a very long time. It is an entirely private monetary system that runs itself and does not depend on trust in any central authority to honor its promises. Instead, it relies on trust in the Bitcoin community or network that verifies transactions and maintains the integrity of the system. This system of distributed trust creates bitcoins and produces an automatic, tamper-proof bitcoin money supply process.¹ As such, it avoids the dangers of discretionary monetary policy—namely, quantitative easing, manipulated interest rates, and the need to rely on wise men or women to withstand political pressure or successfully forecast the future. Indeed, under Bitcoin there is no monetary policy at all. There is just an automatic monetary rule dictated by the Bitcoin protocol designed in 2009 by an anonymous programmer using the alias Satoshi Nakamoto.

Cato Journal, Vol. 35, No. 2 (Spring/Summer 2015). Copyright © Cato Institute. All rights reserved.

Kevin Dowd is Professor of Finance and Economics at Durham University in the United Kingdom and a partner at Cobden Partners. Martin Hutchinson is a journalist and author of the Bear's Lair column (www.tbwns.com/category/the-bears-lair). The authors thank Ferdinando Ametrano, Gavin Andresen, Raadhiyah Anees, Steve Baker MP, Roger Brown, Dave Campbell, Akin Fernandez, Dominic Frisby, Jim Harper, Doug Jackson, Gordon Kerr, Jim Rapp, Eric Samieski, Lawrence H. White and Basil Zafiriou for much helpful feedback. We note that several of our readers have expressed serious reservations about our analysis and conclusions. Any remaining mistakes are the authors' own.

¹By convention, Bitcoin with an upper case “B” refers to the Bitcoin protocol that sends and receives payment information; and bitcoin with a lower case “b” refers to the corresponding unit of money.

Bitcoin has been widely hailed as a success and has won a substantial following. Unfortunately, the underlying economics of Bitcoin mean that it is unsustainable and in all likelihood will be remembered as a failed experiment—at best a pointer to some superior successor.

A first-pass intuition into Bitcoin can be obtained from a comparison with the stone money in Milton Friedman's (1992) case study, "The Island of Stone Money." In this story, the people of the island of Yap in Micronesia used as money large round limestone disks transported from the nearby island of Palau. These were too heavy to conveniently move around, so they were placed in prominent places. When ownership was to be transferred (e.g., as part of a dowry, inheritance, or ransom payment), the current owner would publicly announce the change in ownership but the stone would typically remain where it was and the islanders would maintain a collective memory of the ownership history of the stones. This collective memory ensured that there was no dispute over who owned which stones. Similarly, in Bitcoin, the record of all transactions, the "blockchain," is also public knowledge and is regarded as the definitive record of who owns which bitcoins. Both the stone money and Bitcoin share a critical feature that is highly unusual for a monetary system: both systems operate via a decentralized collective memory.

On February 11, 2009, Nakamoto gave an explanation of the thinking behind Bitcoin in an e-mail announcing its launch: "The root problem with conventional currency is all the trust that is required to make it work. The central bank must be trusted not to debase the currency, but the history of fiat currencies is full of breaches of that trust. . . . With e-currency based on cryptographic proof, without the need to trust a third-party middleman, money can be secure and transactions complete." Cryptocurrencies, however, face the problem of "double-spending." As Nakamoto notes, "Any owner could try to re-spend an already spent coin by [digitally] signing it again to another owner. The usual solution is for a trusted company with a central database to check for double-spending, but that just gets back to the trust model. . . . Bitcoin's solution is to use a peer-to-peer network to check for double-spending." Consequently, "the result is a distributed system with no single point of failure."²

²Quoted from <http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source>.

The fact that Bitcoin has no single point of failure is highly significant: it means that it cannot be brought down by knocking out any particular individual or organization.³ It can only be brought down by knocking out the whole network or one of the underlying building blocks on which the network depends.⁴ It can and does operate outside of government control: Bitcoin is a dream come true for anarchists, criminals, and proponents of private money.

Despite its success, the Bitcoin system is unsustainable due to a design flaw at the very heart of the system. The problem is that Bitcoin requires competition on the part of “bitcoin miners” who validate transactions blocks, but this competition is unsustainable in the long run because of economies of scale in the mining industry. Indeed, these economies of scale are so large that the bitcoin mining industry is a natural monopoly. Furthermore, there are signs that competition in this industry is already breaking down. Once that happens, the system will no longer be able to function as it hitherto has. Its key attractions (decentralization, absence of a single point of failure, and anonymity) will disappear; there will no longer be any reason for users to stay with it; and the system will collapse.

How Bitcoin Works

Let’s start by explaining how Bitcoin works.⁵ The first point to appreciate is that the system is based on the use of Public-Key Encryption (PKE) used to digitally authenticate a signature. PKE is the basis of Internet financial security and is widely used to protect sensitive financial information. Each individual user has both a public key, known to everyone, and a private key, known only to it. When Alice transfers a bitcoin to Bob, she adds Bob’s public key to the coin and digitally signs the coin using her private key. When Bob receives

³By contrast, all other successful monetary or payments systems, including private ones, are dependent on a single individual or organization, which then represents a potential point of failure. Two well-known examples are the Liberty Dollar and e-gold, both of which were perfectly respectable (and legal) business enterprises that were highly successful until they were shut down by the U.S. government. For more on these cases, see Dowd (2014) and White (2014).

⁴The two most obvious such pillars are the Internet and the encryption technology on which Bitcoin depends. Trying to bring Bitcoin down is therefore a tall order.

⁵This account of how Bitcoin works draws extensively from Nielsen (2013).

the bitcoin, he is able to verify that only someone with Alice's private key could have sent it to him: he can then be confident that the bitcoin came from her.⁶

There is, however, a problem: how to prevent Alice double-spending the same bitcoin. So let's look at the general problem of Alice making a payment to Bob: there needs to be a transfer and the transfer needs to be validated without allowing Alice to double-spend. With conventional cash, the transfer is easy: Alice hands over a \$1 note to Bob and Bob just needs to check that the note is not a fake. Alice is unable to spend the same dollar note twice: having handed over the note to Bob, she can't then spend it again with Charlie. With a conventional bank check, Alice sends the check to Bob, and the central authority (the bank) verifies the transfer and adjusts Alice's account to prevent her spending the same deposit money twice. Under the Bitcoin system, on the other hand, there is no central authority to validate the transfer and prevent Alice double-spending the same bitcoin. Instead, these tasks are performed by the network itself—that is, by the community of users.

To get the system to perform these tasks we then have to ensure that there is a reliable verification process and we have to incentivize that process. At first sight, the most obvious verification approach would be a voting system: other users could vote on whether the transfer was valid. If the majority agreed, then the transfer would be deemed valid. Unfortunately, this approach does not work. Remember that the users are anonymous and all we really have is user IDs. So Alice can set up a billion different IDs and take over the system: she has one bitcoin, which she then spends with Bob, Charlie, David, and so on, and she uses her billion votes to approve each transaction; she sock-puppets the system by overwhelming it with spam votes, known in the trade as a "Sybil attack." The way round this problem is to make each "vote" costly, and this costliness is achieved by imposing a proof-of-work requirement. In the Bitcoin system, the proof of work is to demonstrate that the user/voter has expended valuable CPU power to solve a difficult mathematical problem. To oversimplify slightly, the veracity of any transaction is then determined by the majority of CPU-adjusted votes. It is now no longer economic for Alice to spam the system.

⁶The principles of PKE were first publicly set out by Diffie and Hellman (1976) and the first PKE algorithm was set out by Rivest, Shamir, and Adleman (1978).

In addition to solving the double-spending problem, Bitcoin must address the incentive problem—namely, the need to incentivize the network to validate any transaction. Users need to be given some reward for spending their valuable CPU power to validate other people's transactions. Under the Bitcoin system, this reward comes from a combination of newly minted bitcoins and a transaction fee. Leaving aside the fee, the analogy here is with a gold miner mining for gold and occasionally finding it. Similarly, the Bitcoin validator goes searching for bitcoins and occasionally gets lucky. For this reason, the process of validating the bitcoin transactions blocks is usually referred to as "bitcoin mining." In the very early Bitcoin system, a successful bitcoin miner was rewarded with 50 bitcoins for each block validated. However, after every 210,000 validated blocks the rewards halve. This halving has occurred once so far, in November 2012, so the current reward for validating a block is 25 bitcoins. This halving will continue roughly every four years so the production of new bitcoins over time will asymptotically go to zero. It is anticipated that the rate of bitcoin production per block will fall to less than 1 satoshi, or 0.00000001 BTC, by 2141. By that point, the total amount of bitcoin produced will be just short of 21 million. Note, too, that because the time to validate each block is fairly predictable (i.e., about 10 minutes), and even this randomness tends to cancel out over the long run thanks to the law of large numbers, then the rate of bitcoin production, and the hence the total mined by any future time, is highly predictable.

The actual amount of bitcoin available to use, however, is unknown, because of the risk of permanent loss. Loss of the private key to a bitcoin wallet results in the permanent loss of the bitcoins in that wallet, for example. Bitcoins can also be lost when hard drives fail or become infected and the user has not backed them up. One early bitcoiner was reported to have had three copies of his bitcoin wallet, but inadvertently managed to erase two of them and then lost his password for the third, in the process losing \$140,000 worth of bitcoin value. There are probably many others like him. Bitcoins can also be lost due to sundry other forms of human error. In one widely publicized case (reported, e.g., in ITN 2013), an IT worker in Wales generated 7,500 bitcoins on a website in 2009. He then forgot all about them and later threw out his hard drive in a clear-out. By the time he realized his mistake, his bitcoins were worth about £4.6 million but were still on his

hard drive, which was now buried in a landfill in Newport. He was now a virtual millionaire in more ways than one. He went over to the site, which was now the size of a football field, and was told that his hard drive was probably buried 3–5 feet beneath the surface somewhere on the site. He subsequently spent a few weeks digging around for it, but to no avail, and the local council has since barred anyone else from looking for it. There are no hard estimates of how many bitcoins have been permanently lost in such ways.

An important principle of the Bitcoin system is that it does not assume that all miners are honest when validating transactions. To deal with possible dishonesty, it seeks to establish incentive-compatibility to reward honest miners and make dishonest mining unprofitable. To understand how this works, we can think of the mining process as a competition to approve transactions, and the cost of entry to this competition is a certain amount of CPU power. A miner's chance of success in the next competition is then roughly equal to the proportion of total computing power that it controls, so a miner with 1 percent of the computing power being used to validate transactions has a roughly 1 percent chance of winning the competition. It then follows that, provided a lot of computing power is being brought to bear on the competition, and provided that most miners are honest, a dishonest miner is likely to have only a small chance to corrupt the validation process, unless it expends a huge amount of computing resources—and that will be costly. The idea (hope?) is then that dishonest mining will rarely occur, for the simple reason that it is not a profitable strategy.

To spell out the argument further, we need to consider how the blockchain works. For the system to work properly, we need the network to agree on the order in which transactions occurred, as otherwise it may not be clear who owns which bitcoins. To achieve this objective, each block in the blockchain includes a pointer to the previous block (in fact, this pointer is just a hash of the previous block) until we go back to the original block, the Genesis block, and we end up with a linear blockchain.

Sometimes, however, a fork will appear on the chain. This might happen if two miners happen to validate different blocks at almost the same time, both broadcast their newly validated blocks to the network, and some network members update the blockchain using one block, while others update it using the other block. This fork can

cause confusion, as it is no longer clear which set of transactions is to be regarded as valid. To get around this problem, the Bitcoin protocol stipulates that if a fork occurs, miners should keep track of both forks, but at any point in time, they should only work on the longer one. This rule means that once one fork gets a sufficient lead over the other, then it will become the generally accepted one and the blockchain will revert to the linear form.⁷ However, any pending legitimate transactions in the abandoned fork will eventually be validated in the successful fork, so all legitimate transactions will still be validated regardless of the fork in which they first appeared.

One can then imagine Alice attempting a variety of different double-spend attacks. One is for her to attempt to validate a block that includes a first spend with Bob and then a double-spend with Charlie. But even if she gets lucky and is able to validate her block—and the odds are that she won't—the attack won't work because other users will refuse to confirm her validation. A second possible attack is to spend the same bitcoin with Bob and Charlie, but broadcast each transaction to different parts of the network in the hope that both transactions will be confirmed. However, this won't work either, as the network will eventually confirm only one transaction. A third strategy is to spend a bitcoin with Charlie, wait for the transaction to be fully confirmed, and then pretend to be Bob and spend it on him. To do this, Alice creates a new fork, but her new fork is already six confirmations behind the leading fork. She then faces the problem of how to overtake the leading fork, but to do that she would have to win the competition to validate at least the next six blocks. If she has 1 percent of the network's computing power, the probability of her winning the next seven block validations is then $1/100^7 = 0.00000000000001$ percent. So her chances of success are vanishingly small—she may as well be honest instead.

This discussion is neither rigorous nor exhaustive, but it does illustrate how the Bitcoin system is designed to counter dishonest behavior by being incentive-compatible. How well it achieves this objective is however another matter, and we shall return to this issue presently.

⁷In fact, under the Bitcoin protocol, a transaction is not considered confirmed unless it is part of the longest fork and at least five later blocks follow it in that fork, i.e., six confirmations are required before a transaction is regarded as finalized.

Why Bitcoin Mining Is a Natural Monopoly

As we have seen, a central feature of the Bitcoin system is that it relies on competition between independent miners to maintain the integrity of the blockchain: it relies on a form of perfect competition between atomistic miners. Unfortunately, this perfect competition is not a stable equilibrium, because the atomistic miners have an incentive to collude, or at least to merge.

The Incentive to Merge

Suppose that one bitcoin is to be mined over the next 10 minutes. Now consider any two miners: If they work separately, each gets a certain expected return and there is a certain probability of it successfully mining the bitcoin in the next block. If they merge their operations into a mining pool, they have the same expected return per miner but a higher probability of obtaining that return.

Assume that there are n independent miners, each of whom faces a probability p of mining the next bitcoin. For any individual miner, the expected return from mining the next block is p bitcoin, since he has a probability p of a return of 1 bitcoin. Now suppose that two miners form a pool. The mining pool's expected return is now $2p$ bitcoin, because it has a probability $2p$ of mining the next bitcoin. The expected return for each pool member is therefore p bitcoin. So by joining the pool, the individual miner gets the same expected return as it would get if it mined on its own, but it doubles its probability of getting a positive return. Assuming miners are risk-averse, it then makes sense for any two miners to pool their resources and share the profits.

However, if it makes sense for any two miners to form a pool, it also makes sense for any group of miners to form a pool. Thus, the original competition between individual miners in the Bitcoin system consolidates into competition between ever growing mining pools: perfect competition gives way to oligopoly.

Now suppose that the process of consolidation has produced four mining pools of equal size. Then each pool gets an expected return of 0.25 bitcoin because the probability that any one pool successfully mines the next bitcoin is 0.25. Each of the $p/4$ pool members then has an expected return of p bitcoin, being a 0.25 chance of a $4p$ return. The original individual miners are clearly better off in the bigger pools than they were operating independently, or even in smaller pools.

But the best outcome is simply for all the remaining pools to merge into one super pool. The expected return to that super pool is 1 bitcoin and it would achieve that expected return with a probability of 1. (This is an oversimplification because we ignore the possibility of new entrants coming into the mining market, but this is just a detail.) The expected return to each pool member is again p bitcoin, which it achieves with a probability approaching 1.

We then arrive at the conclusion that the production of bitcoin is a natural monopoly, an industry in which it is technically more efficient to have one producer rather than many. So even if there are a large number of producers/miners to begin with, there will be economic pressure on them to combine.

The Negative Externalities of Competitive Mining

There is also a second reason why bitcoin mining is a natural monopoly: the negative externalities of competitive mining. The expected marginal revenue for an individual miner increases with the amount of CPU power it devotes to bitcoin mining, but the difficulty of the problem it must solve (and hence the marginal cost involved) increases with the amount of CPU power expended across the entire network. The result is that individual miners do not take into account the negative cost externalities that their own activities impose on other miners. The competitive equilibrium is where the expected marginal private benefit (MPB) from mining is equal to the expected marginal private cost (MPC) of mining, but since the latter is less than the marginal social cost (MSC) of mining, we get an equilibrium in which excessive resources are devoted to mining-related activities. In particular, there is excessive use of energy and excessive investment in computing resources.

There is also an externality on the benefit side as well. An increase in mining activity will produce a positive expected MPB to the miner concerned, but the marginal social benefit (MSB) must always be zero, because the Bitcoin protocol dictates that there are only so many bitcoins to be had. So even if the activity were costless, seeking the private benefits of bitcoin mining is a zero sum game—that is, one miner's gain must be another miner's loss. Moreover, from the perspective of social efficiency: since the MSB is zero, the socially

optimal MSC is the lowest possible, which is a far cry from what the competitive mining equilibrium delivers.

Consider the evidence. The daily revenue to bitcoin miners since January 2009 shows that this revenue was initially very low, with occasional peaks, the biggest being just over \$5 million in late 2013. Miners' revenue then fell erratically and was \$1,331,071 as of December 13, 2014, equivalent to \$370 per bitcoin mined. The price of bitcoin the same day was \$349. The difference between the revenue per bitcoin and the price of bitcoin reflects the average transaction fee that day.⁸

We also need to consider the cost of production, and these have risen enormously. In the very early days of bitcoin, a typical home PC could mine hundreds of bitcoins a day, but by late 2014, PCs have long since been obsolete for bitcoin mining, and a state of the art mining computer, an Application Specific Integrated Circuit (ASIC), can expect to mine only fractions of a bitcoin a day. A key component of the cost of bitcoin mining is the "hash rate"—that is, the number of gigahashes (billions of hashes) calculated by the network each second. In January 2010, the hash rate was 0.01 or 1 percent but is now more than 300 million, increasing by a factor of more than 30 billion. We also have to consider that as the rate of production of bitcoins halved in November 2012, the hash rate per bitcoin actually mined increased by a factor of over 60 billion, as of December 13, 2014.⁹

However, what matters from a cost perspective is not the hash rate as such, but the cost of producing it, and this cost has been falling sharply over time. A good rule of thumb here is Koomey's Law, which suggests that the amount of battery power needed to achieve a fixed computing load has been halving about every 18 months since the 1950s (Greene 2011). If we count five years since January 2010, the amount of energy per hash will have fallen by a factor of between 4 and 5 over this period, and the cost-adjusted hash rate will then have risen by a factor of around only 10 billion or so. This is, needless to say, still an enormous rate of growth.

There is also the question of the profitability of bitcoin mining. Some insight into this issue is provided by a recent study by Hass McCook (2014). To start with, we need to keep in mind that the

⁸Downloaded from blockchain.info on December 13, 2014.

⁹*Ibid.*

profitability of mining depends in part on the local cost of electrical power and the cost/efficiency of the mining machine used, and these vary considerably. He points out that the rate of increase of the difficulty of the underlying computing problem, which is reflected in the rate of increase of the hash rate, is typically between 10 and 20 percent a fortnight, implying that the useful life of most mining equipment is only about 3 to 6 months. Putting these points together, he comes up with some ballpark figures. He suggests an average capital cost of \$530 for each mined bitcoin, and a corresponding average operational cost of about \$67 per bitcoin, meaning a total cost per bitcoin of just under \$600. By a curious coincidence, this was about the same as the price of a bitcoin at that time. Of course, we should keep in mind that this is very much a ballpark average figure and there will be considerable variation around it. However, if we give it some credence, we can say that by July 2014, the average profitability of bitcoin mining had declined to the point of marginality. Since then, the price of bitcoin has fallen by just over 40 percent, and the hash rate—and presumably the cost of mining—have risen further. We can then reasonably infer that most bitcoin mining is now unprofitable.

We should now expect to see miners exiting the field by letting their computers become obsolete without replacing them. The CPU power devoted to mining would then fall to the point where it became profitable to resume mining again, and the cycle would repeat itself. However, we should keep in mind that the profitability of mining also depends on the price of bitcoin. If the price of bitcoin increased, then currently unprofitable mining operations would become profitable and more CPU power would be devoted to mining again. However, if the price of bitcoin fell, currently profitable operations would become unprofitable, and there would be further downward pressure on mining CPU power. In that case, mining capacity would keep falling until it reached the point where mining again became profitable—or, alternatively, if the price of bitcoin kept falling, mining capacity would fall indefinitely.

It is also insightful here to consider the incentives faced by a bitcoin miner. For the sake of illustration, let's take as given McCook's cost estimates just set out. If the price of bitcoin is higher than about \$600 and expected to remain so, then bitcoin mining is profitable. We would then expect the miner not just to mine for bitcoin, but also to mine on the maximum scale it can manage, investing all it can in further mining. This incentive to scale up helps explain why mining

capacity has expanded so much, and is a critical and destabilizing feature of bitcoin market dynamics. If the bitcoin price is below \$600, but expected to stay above \$67, then it is no longer economic to continue investing in bitcoin mining, because the price obtained no longer covers both the capital and operational costs involved. However, the marginal costs of bitcoin mining are still covered, so the miner will use existing capacity to continue mining, but will not maintain that capacity. Mining capacity will then rapidly run down and effectively disappear in six months or so. Lastly, if the price of bitcoin were to fall below \$67, then mining activity would no longer cover even the operational costs, and the miner would stop mining altogether.

Returning to our natural monopoly theme, it is plainly obvious that the industry could reap considerable benefits if the major players were able to form a cartel or unite into a single outfit. This outfit could stop further investment in bitcoin computing and switch off most of the existing mining computers. Thereafter, it would maintain the blockchain using a much smaller number of computers—perhaps just a couple of servers—expending much less CPU power. The cost gains achieved in this way could then be shared around the members as additional profit. The potential savings are considerable: If we take McCook's figures for July 2014, then most of the nearly \$600 mining cost for each bitcoin could be avoided. With $24 \times 6 \times 25 = 3,600$ bitcoins currently being produced each day, there would be a cost saving of the greater part of $3,600 \times \$600$ or just over \$2 million per day.

Implications of a Bitcoin Natural Monopoly

These tendencies to centralization are totally destructive of the Bitcoin system. The central innovations of Bitcoin are distributed trust and the absence of any single point of failure. The system has worked because users could trust the network as a whole to maintain the integrity of the system, and so avoid the need to trust any individual body such as a central manager; and the system couldn't be brought down, because it had no single point of failure. However, this model only works so long as the network consists of a sufficient number of competing mining entities. Once the individual miners coalesce into a dominant player or a group of big players that are able to form a cartel, then that dominant

player/cartel has control over the system—it decides which transactions are to be deemed valid, and which are not. We then have to trust that entity not to abuse its position and are back to the trust model, not a self-regulating one. That same entity also becomes a point of failure for the system as a whole. In short, the fundamental contradiction in the Bitcoin system is that it requires mining to be competitive, but it creates a mining market structure that is a natural monopoly, and the two are incompatible because of the latter's tendencies toward centralization.

Going back to our island of stone money, it is as if everyone woke up one morning unable from that point on to remember who owned which stones. However, one individual still claims that he can remember and helpfully offers to remember for everyone else. One wonders how well that would work.

In each case, we have a system that depends on collective memory to operate, so it is difficult, to say the least, to see how it could survive the undermining of the collective memory on which it depends—especially when there is no instantly credible replacement to the original collective memory trust model that has just been eviscerated.

There are also further casualties to the Bitcoin system. Once a dominant player arises, it cannot possibly operate in a clandestine fashion beyond the knowledge of law enforcement; it cannot operate anonymously like Silk Road on a dark web browser and thumb its nose at the government. If it cannot operate anonymously, then it cannot escape regulation by the state and would have to submit to whatever state law requires in terms of tracking individual users and reporting results to legal authorities. The combination of the Bitcoin system being controlled by a publicly known entity, the requirement that that entity submit to the law, and evolving de-anonymization technology would then be fatal for user anonymity.¹⁰ Users of bitcoin

¹⁰It often said that Bitcoin is anonymous, but the truth is that it is actually pseudonymous and operates via user IDs. The anonymity (or otherwise) of Bitcoin then depends on whether a user's true identity can be inferred from these user IDs or the blockchain record. Even back in 2011, Jeff Garzik, a member of the Bitcoin dev team, was openly cautioning bitcoiners not to take Bitcoin anonymity for granted. He warned that law enforcement could parse the transactions flow to track down users in the same way that they can detect suspicious money flows. Recent studies such as Merklejohn et al. (2013), Möser (2013), and Biryukov, Khovratovich, and Pustogarov (2014) look at various ways in which anonymity can be broken and confirm that most bitcoin transactions are not truly anonymous (see also Nielsen 2013).

could no longer operate on the basis of any assumed anonymity, and those who unwisely carried on using bitcoin for illegal purposes would soon fall foul of law enforcement. Anonymity on part of miners or bitcoin users would then disappear. Indeed, the likelihood is that the government would destroy anonymity at a stroke by requiring that the bitcoin dominant player insist that any users of the network openly register themselves by providing photo ID, Social Security numbers, and proof of address in much the same way that bank depositors are required to provide the same information to their banks for anti-money laundering and similar purposes. The demand for bitcoin to make any illicit transactions (e.g., to buy illegal drugs on Silk Road and similar sites, or to move funds around illegally) would then disappear.

One has to ask why users of bitcoin would rationally continue to have any confidence in the Bitcoin system when it becomes apparent that its key attractions—decentralized trust, no single point of failure, and the possibility of using it for anonymous transactions—had all been compromised. Remember also that the willingness of any individual to accept bitcoin is entirely dependent on his or her confidence that other people will continue to accept it. There is nothing in the system to anchor the value of bitcoins because, unlike gold or tulips, bitcoins have no alternative use value.

Nor is there any rational reason to trust in the dominant player/cartel to behave itself. Trust comes from credible assurances—it comes from credible precommitment, a willingness to post performance bonds, and to submit to account—and there is no way that a shadowy dominant mining pool can provide such assurances. In any case, there is no reason to want to trust such an entity when you can use safe and reputable systems such as PayPal or the better banks.

The whole Bitcoin system then becomes a house of cards. There is nothing within the system to maintain confidence in the system, and anything—a scandal, a government or regulatory attack, or any factor that triggers a loss of confidence—could lead to a run that brings down the entire system. It is then rational to sell out before that happens, and if enough individuals think this way, their expectations will become a self-fulfilling prophecy. There will be a stampede for the exit; the price of bitcoin will drop to bitcoin's intrinsic value, zero; and the system will collapse.

The Emergence of the Big Mining Pools

Indeed, it may not take the emergence of a monopolist to bring on the crisis, but merely the emergence of a mining pool big enough to threaten the system. If we look at the bitcoin mining market, we see that mining pools have been growing rapidly and are already very large. An idea of their size can be seen from the distribution of hash rate power across the bigger mining pools. As of December 13, 2014, Discus Fish had 25 percent of the hash rate distribution and GHash.IO had 17 percent. The other bigger pools were somewhat smaller, but 20 percent of the distribution is “unknown.”¹¹ It is therefore conceivable that some of this unaccounted hash rate is due to the other big players. Thus, the shares of the big players might be larger than they seem to be. The distribution also moves around a reasonable amount from day to day, so the data should be regarded as a snapshot from a volatile motion picture. Nevertheless, it indicates that some pools are now large enough to be a source of threat to the system as a whole.

The most interesting of these mining pools is the secretive GHash.IO. Its business model is based on zero fees and cloud hashing, in which miners rent out hash power from its sister outfit, CEX.IO, which supplies the hardware, hosting, and maintenance. Its webpage offers instant payouts, zero fees, 24/7 support, and SMS messaging. We are told that GHash was launched in July 2013, is trusted by 300k users, and (with about one third of the network's hashing power) is already the largest mining pool in the Bitcoin community. However, the webpage offers no phone numbers, offers no credible reassurances, and has no audited accounts that we know of. It also says nothing about the people involved or where it is based, although further investigation suggests that the person behind GHash is one Jeffrey Smith. The photo on his Twitter page (Jeffrey Smith@jeff_smith01) shows him as a gentleman in his early 20s wearing a T-shirt, but even so, we still cannot be sure whether this is the person behind GHash or Mr. Smith himself. In short, we have no real idea who is behind the allegedly biggest mining pool in the Bitcoin community. However, we can safely assume that he is probably not some latter day John Pierpont Morgan, the premier

¹¹Downloaded from blockchain.info on December 13, 2014.

financier of his age, operating out in the open, putting his considerable wealth and reputation on the line with each deal he makes and even choosing to operate under unlimited liability to give his counterparties credible reassurance.

Once mining pools reach a big enough size, there is a danger that one of them will launch a 51 percent attack in which a mining pool achieves 51 percent or more of the mining power on the network and then uses this power for malicious ends that the system was designed to prevent. It can keep all the mined bitcoins to itself, double-spend bitcoins, selectively reject competing miners' transactions, extort high fees from those with large holdings by making certain addresses unspendable, or launch a denial-of-service attack against the network. Another threat is "selfish mining"—that is, a strategy in which a pool

keeps its discovered blocks private, thereby intentionally forking the chain. . . . Rational miners will preferentially join selfish miners to reap the higher returns. . . . Such a selfish mining pool will quickly grow to become a majority, at which point the pool will be the only creator of blocks, the decentralized nature of the currency will collapse, and a single entity, the selfish pool manager, will control the system [Eyal and Gün Sirer 2013b: 2].

The possibility of a selfish mining attack is also significant because it proves that the Bitcoin system is not fully incentive compatible. Eyal and Gün Sirer offer various fixes to this threat. But even in their best-case scenario with their fixes in place, their analysis shows that the system can be taken over by a mining pool that has 33 percent of the network's hash power. Even in their best-case, we would need at least two-thirds of the miners to remain honest, not 50 percent as is widely believed, for the system to be safe—and without any fixes it is not safe at all. Their conclusions are stark:

Bitcoin is broken. And not just superficially so, but fundamentally, at the core protocol level. We're not talking about a simple buffer overflow here, or even a badly designed API that can be easily patched; instead, the problem is intrinsic to the entire way Bitcoin works. All other cryptocurrencies and schemes based on the same Bitcoin idea, including Litecoin,

Namecoin, and any of the other few dozen Bitcoin-inspired currencies, are broken as well. (Eyal and Gün Sirer 2013a)¹²

The Bitcoin community had long been aware of threats such as 51 percent attacks. However, the conventional view among experts was that they were not as serious as they appeared to be. It would be obvious from the blockchain that an attack was under way, an attacker wouldn't be able to get its proceeds out of the Bitcoin system and there were various ways to deal with an attack in the unlikely event it did occur (e.g., the Bitcoin dev team would fix the problem or the community would pressure a miscreant to behave responsibly). One commentator even went so far as to describe the threat of a 51 percent attack as a boogeyman. The argument was also repeatedly made that it would not be in any pool's own interest to achieve a 51 percent position, not least because this would be to invite mutually assured destruction. Bitcoin insiders also pointed to the community's success in dealing with previous cases where a mining pool had become uncomfortably large.

A case in point occurred in April 2013, when the mining pool BTC Guild obtained a large portion of the network hashrate and provoked concern over its market power. In response, BTC Guild adopted a mitigation plan that saw its market share fall again. This solution worked thanks to the commitment of those involved to the Bitcoin "rules of the game" and their shared understanding of the dangers posed by the emergence of large mining pools. When a pool became too big not only would miners voluntarily leave the pools of their own choice, but the mining pools themselves would voluntarily reverse their own growth by raising fees and halting new registrations. A big pool would then reduce in size and the threat would recede.

¹²It is important to stress that there are alt currencies that are not based on Bitcoin-like mining protocols. One of these is Ripple, which ranks second to Bitcoin in terms of market cap: it currently has just over 9 percent of the cryptocurrency market share, as opposed to Bitcoin's 84 percent. Ripple has a pre-mined currency supply and validates blocks through a consensus system based on trusted validators, not proof-of-work "mining" as under Bitcoin. Its protocol deters collusion and the fact that validators are known entities deters any Sybil attack. Another is BlackCoin, which ranks #24 in the cryptocurrency league with a market share of 0.02 percent. In this system, validation is based on "proof-of-stake": a node that generates a block has to prove that it has access to a certain amount of coins before being accepted by the network. Generating a block involves sending coins to itself, which then proves the ownership and hence the stake in the system.

However, arguments to the effect that no one would engage in behavior that threatened the system, because everyone has a stake in it, are based on simplistic assumptions about the self-interests of the parties concerned. Take an individual miner or small mining pool. It is often stated that such outfits will have a commitment to the system and take a long-term view because of their investment in mining equipment. Yet this argument overlooks the point that ASICs have a useful lifetime of only months. Their only incentive is to extract as much value as possible before their investment becomes obsolete: this is not a long-term investment but its opposite, pure short-termism. Competition between mining pools is also a complicated affair with enormous scope for game-playing, especially with the rapidly evolving technology involved. A big player will also have market power, and we should expect that it will use that power to its own advantage. Like a big bank or utility company, it can practice price discrimination, it can exploit behavioral traits in its customer base, and it can exploit customers with deliberately complicated opaque charge regimes. A dominant miner can also

engage in subtle attacks that are hard to detect. [One of these is] Transaction Differentiation [in which a] 51 percent miner can simply render certain Bitcoin addresses (what clients perceive as “wallets”) either unspendable or highly deprioritized unless a high mining fee is paid. This is tantamount to ransom. In effect, the miner would turn to the Winklevii, who have large Bitcoin holdings, and say “my, my, my, nice fat wallet you’ve got there, you’ll have to attach a 1 percent mining fee if you want to ever spend those coins again” while brandishing the virtual equivalent of a steel pipe [Eyal and Gün Sirer 2014].¹³

¹³The Bitcoin system is also open to a large variety of other attacks. These include: (1) block withholding attacks, in which a pool infiltrates another, but discards any full proofs-of-work and provides only partial proofs-of-payment to its victim, which is thus deceived into thinking that the attacker is doing effective mining; the attacker is then rewarded for its apparent effort when in fact it never contributes to the victim’s revenue stream (Eyal 2014); (2) DOS attacks on the most important nodes in the network, which might seriously disrupt it; (3) side-channel attacks which exploit weaknesses in crypto security to extract information about private keys and enable the attacker to steal bitcoins (see Bengier et al. 2014); (4) any number of malware attacks; (5) the usual market manipulation tactics; and, as noted elsewhere, (6) the threat of attack from governments.

Then along came GHash.IO, which broke the mould in two different ways. First, it had an innovative business model that made it very attractive to individual miners. Apart from the obvious attraction of zero fees, cloud hashing and hosted mining meant that miners no longer had to bear the inconvenience and risks of having their own mining equipment. Not surprisingly, GHash's market share then rose rapidly and it very soon became the biggest mining pool. The second distinctive feature of GHash was its attitude: it wasn't prepared to play by the rules of the game.

In January 2014, GHash.IO achieved 42 percent of the network's hash rate and there was renewed concern about the threat this posed to the network. In response, GHash.IO reassured the community that it would "take all necessary precautions to prevent reaching 51 percent of all hashing power" to maintain the stability of the network.¹⁴ However, GHash pointedly refused to move from the zero fee model and didn't deliver on its promises to keep its market share down. The issue then flared up again on June 15, 2014, when it was reported that GHash had repeatedly contributed more than 51 percent of the network's hashing output for periods as long as 12 hours. As Gün Sirer observed, "having a single entity in GHash's position, of holding 51 percent of the mining power, of being in a monopoly position, of being able to launch any of these attacks at will, completely violates the spirit and intent of Bitcoin as a currency."¹⁵ The very next day, June 16, GHash rushed out a statement to reassure everyone of its continued commitment to the Bitcoin system:

Our investment, participation and highly motivated staff confirm it is our intention to help protect and grow the broad acceptance of bitcoin and categorically in no way harm or damage it. We never have and never will participate in any 51 percent attack or double spend against bitcoin.¹⁶

¹⁴These precautions included a mitigation plan that included allowing CEX.IO customers to point their hashing power toward other mining pools besides GHash. This was a non-solution if ever there was one: the hashing power so diverted would still be under the control of GHash.io/CEX.io and its hashing power would merely *appear* to diminish. Such a "solution" makes the underlying problem worse by hiding it.

¹⁵<http://arstechnica.com/security/2014/06/bitcoin-security-guarantee-shattered-by-anonymous-miner-with-51-network-power>.

¹⁶<http://www.coindesk.com/ghash-io-never-launch-51-attack>.

This statement will, we believe, ultimately prove to be bitcoin's death knell. Bitcoin, the ultimate system designed to avoid the need for trust, is now reduced to relying on trust in the dominant mining pool not to attack it.

In mid-July, Bloomberg journalist Leonid Bershidsky (2014) wrote:

Trust will kill bitcoin, the stability of [Bitcoin] and all the infrastructure that has grown around it in the past five years . . . now depends on the goodwill of a few people whose names nobody knows. "Jeffrey Smith"—the name used by GHash.IO's only spokesman—is likely a pseudonym.

There's no guarantee that a certain group of people doesn't control more than 51 percent of Bitcoin emission even now: GHash.io is currently 39 percent, but another 17 percent is ascribed to "unknown" by Blockchain.info, the best source on mining pool shares. . . .

I'm not willing . . . to ignore the risk that a group of people who won't even tell me who they are might take the money and run.

Trusting them would be akin to acknowledging them as the system's central bank, something Nakamoto emphatically didn't want it to have. Paradoxically, it's easier to trust central banks that print fiat money, because the people in charge and their interests and motives are more or less known, or at least are the subject of much study and speculation.

However, the problem wasn't just that GHash.IO was now placed to compromise the system, again. Part of the problem was its attitude. It still claimed to share the community's concern that mining should be decentralized, but it had done nothing to promote decentralization despite its earlier promises to scale itself back. It continued to refuse to raise its fees: higher fees were only a "temporary" solution and GHash.IO was looking for something more permanent, it said. It then painted itself as the victim: "You cannot blame GHash.IO for being the #1 mining pool," Smith told an interviewer, as if it being the #1 pool had been outside of its control. It had blatantly gone past the 51 percent barrier knowing very well its psychological impact: it was openly flaunting its power and didn't seem to care what the others thought. It then declared its intention to host a round table with the other leading pools and the Bitcoin

Foundation to find ways to decentralize the mining industry, simultaneously declaring its adamant opposition to the traditional solutions of high-level Bitcoin idealism—that is, self-restraint, to which the other mining pools (apparently) still adhered.

Yet the underlying problem is not GHash’s attitude or any secret agenda, or the fact that the community lacks an effective means of bringing “offenders” under control.¹⁷ Instead, the problem is that many of the proposed solutions to the problems posed by large mining pools are not incentive-compatible.

One such “solution” is to suggest that individual miners would (or should) behave responsibly because of *their* stake in the system. After the GHash.IO threat surfaced in January 2014, one commentator issued a stark statement:

This is a WARNING to all Bitcoin Miners. We sincerely request that miners using GHash.io . . . leave the pool for other smaller pools to create a more decentralized Bitcoin hash distribution.

From the perspective of the community, it would have been a good thing for individual miners to heed this warning and switch away from GHash.IO. But why should they? Yes, very large mining pools are antisocial, but why should any individual miner forgo his own self-interest for the greater good? One has to bear in mind that for an individual bitcoin miner, a big pool like GHash is an attractive proposition—it charges no fees, offers a fairly stable return because of its size, and even takes away the hassle and risk of buying and maintaining one’s own ASIC miner. Why should such a miner forgo those benefits when so many other miners are quite happy to enjoy them? Again, one is left with only an appeal to put conscience over self-interest. In any case, one can no longer assume that individual miners even have much of a stake in the Bitcoin system. In the past, there were major setup costs to mining, and most

¹⁷ Their only weapons are bad publicity, ostracism, and DoS attacks, all of which have been tried on GHash.IO without much success. Bitcoiner critics are also constrained in that they can’t be too aggressive in their criticisms without the risk of undermining confidence in Bitcoin itself. The truth is that if one big player refuses to play by the “rules of the game,” there is nothing much that the Bitcoin community can do about it.

miners had some commitment to the system, ideological as well as intellectual. However, as mining has become mainstream and much easier, you can now become a bitcoin miner without knowing much about Bitcoin or even giving a damn about it: a friend tells you that Bitcoin mining is a good prospect, so you take a few minutes to register at GHash, pay your money, and wait for the bitcoins to roll in.

“Is this really Armageddon?” asked Eyal and Gün Sirer (2014). “Yes, it is,” they answered. We agree. As they continued:

The fact is, this [GHash positioning itself to threaten the network] is a monumental event. The Bitcoin narrative, based on decentralization and distributed trust, is no more. True, the Bitcoin economy is about as healthy as it was yesterday, and the Bitcoin price will likely remain afloat for a while. But the Bitcoin economy and price are trailing indicators. The core pillar of the Bitcoin value equation has collapsed.¹⁸

Conclusion

The fact that the Bitcoin system was designed not to require trust, but now depends on it, suggests that it is living on borrowed time. Human nature being what it is, the temptation to abuse that trust will prevail. The history of central banking is full of breaches of such trust, and it would be perverse to assume that the big

¹⁸ We should add that Eyal and Gün Sirer themselves do not draw our conclusion that bitcoin is doomed. Yes, it is Armageddon, they say, but they go on to state that there is life after Armageddon and much of their work is devoted to suggesting fixes of one sort or another that they claim would alleviate or solve some of the problems posed by game-playing mining pools. We would claim that our conclusion is the simplest and most natural one, however, and that their fixes, *if* adopted, would only delay the inevitable and may not be adopted anyway. In fact, there has been a great deal of discussion on the blogosphere about the large mining pool problem. Most proposed solutions boil down to encouraging decentralization, a good example being the “getblocktemplate” mining protocol developed over mid 2012, the essence of which is to move block creation from the pool operator to the individual miner. We cannot even attempt here to provide any analysis of these proposals, other than to repeat our view that these are at best palliatives that may ameliorate but do not root out the underlying problems posed by centralizing tendencies in the Bitcoin system. We would also add that our message that Bitcoin will bite the dust has received a mixed but mostly negative response from the Bitcoin community.

bitcoin players would be immune to the same human failings as central bankers.

Our best guess is that in the short term there will continue to be a drip-by-drip erosion of confidence as the realization grows that the system is compromised. The current bitcoin oligopoly cannot resolve this problem: like the Roman triumvirates, it is unstable and the principal parties involved cannot agree on a lasting solution; they also lack the ability to provide the necessary credible assurances anyway. Humpty is well and truly broken and neither all the King's horses nor all the King's men can put him back together again: it can only be a matter of time before the whole unsteady edifice will collapse.¹⁹

Even in the unlikely event that it survives into the medium run, we would still rate its longer-term chance of survival as zero. First, we should remember that a recurring theme in the history of innovation is that the pioneers rarely, if ever, survive. This is because early models are always flawed and later entrants are able to learn from the mistakes of their predecessors. There is no reason why Bitcoin should be an exception to this historical rule. The second reason is that in the very long run bitcoin would be uncompetitive against efficient closed-wall systems such as PayPal or COEPTIS, the successor to e-gold. Once the production of bitcoins becomes insignificant, then the Bitcoin system will entirely depend on transaction fees to cover its operational costs, and its fee levels would be higher than those of more traditional payment systems because of the need to maintain excess hashing and excess capacity to deter new entrants into the transaction validation business. Put differently, Bitcoin can never achieve the technical economic efficiency of competitors that can operate with a very small number of servers, or even just one. In the very long term, when there are no new bitcoins being produced to subsidize the validation process, the Bitcoin system will no longer be

¹⁹ Many of the alt cryptocurrencies have similar weaknesses to bitcoin. Lawrence H. White (2015) documents that many of these have already collapsed, including Terracoin, FreicoIn, BBQCoin, Megacoin, Mooncoin, CryptCoin, ScotCoin, Bitgem, and CrtCoin. Of these, Terracoin experienced a decline of 97 percent, Megacoin experienced a decline of over 98 percent, and all the others mentioned experienced declines of over 99 percent. This evidence suggests to us that mining-based alt currencies have short life-spans and soon succumb to inherent tendencies toward centralization, takeover, and collapse.

able to compete.²⁰ Last but not least, there is still the problem that Bitcoin is not backed by anything.

Yet the undeniable achievement of Bitcoin is that it demonstrates the practical *possibility* of fully decentralized monetary systems based on the principle of distributed trust rather than central authority. Like the Wright brothers, it shows that such systems can fly, but it does not demonstrate that they can stay in the air for too long. We would therefore regard Bitcoin as an instructive creative failure, but we are hopeful that the lessons to be drawn from the experience of Bitcoin and other cryptocurrencies will lead to superior private currencies in the future, crypto and otherwise. Further experimentation in the private money space is therefore to be welcomed. As is the nature of capitalist innovation, most of these experiments will doubtless fail, but a few will succeed—perhaps hybrids of crypto and gold. As the old saying goes: “Make new friends, but keep the old. One is silver and the other is gold.”

References

- Benger, N.; van de Pol, J. N.; Smart P.; and Yarom, Y. (2014) “Ooh Aah. . . Just a Little Bit’: A Small Amount of Side Channel Can Go a Long Way.” Mimeo: School of Computer Science, University of Adelaide.
- Bershidsky, L. (2014) “Trust Will Kill Bitcoin.” *BloombergView* (17 July 17): www.bloombergview.com/articles/2014-07-17/trust-will-kill-bitcoin.
- Biryukov, A.; Khovratovich, D.; and Pustogarov, I. (2014) “Deanonymisation of Clients in Bitcoin P2P Network.” Mimeo: University of Luxembourg, arXiv:1405.7418v3 [cs.CR] 5 Jul 2014.
- Diffie, W., and Hellman, M. E. (1976) “New Directions in Cryptography.” *IEEE Transactions on Information Theory*. IT-22 (6): 644–54.
- Dowd, K. (2014) *New Private Monies: A Bit-Part Player?* Institute of Economic Affairs, Cobden Centre Hobart Paper No. 174. London: Institute of Economic Affairs.

²⁰ One should also bear in mind a related problem with Bitcoin’s competitiveness: the fact that a block takes about 10 minutes to validate can be a considerable inconvenience for point-of-sale transactions, especially when there are competing payments systems that can offer almost instantaneous payment verification.

- Eyal, I. (2014) “The Miner’s Dilemma.” Blog posting at <http://hackingdistributed.com/2014/12/03/the-miners-dilemma>.
- Eyal, I., and Gün Sirer, E. (2013a) “Bitcoin Is Broken.” Blog posting at <http://hackingdistributed.com/2013/11/04/bitcoin-is-broken>.
- _____ (2013b) “Majority Is Not Enough: Bitcoin Mining Is Vulnerable.” Available at www.cs.cornell.edu/~ie53/publications/btcProcArXiv.pdf.
- _____ (2014) “How a Mining Monopoly Can Attack Bitcoin.” Blog post at <http://hackingdistributed.com/2014/06/16/how-a-mining-monopoly-can-attack-bitcoin>.
- Friedman, M. (1992) “The Island of Stone Money.” In M. Friedman, *Money Mischief*, 3–7. New York: Harcourt Brace Jovanovich.
- Greene, K. (2011) “A New and Improved Moore’s Law: Under ‘Koomey’s Law,’ It’s Efficiency, Not Power, that Doubles Every Year and a Half.” *Technology Review* (12 September).
- ITN (2013) “Man’s £4m Bitcoin Fortune Lost in Rubbish Dump.” Available at www.itv.com/news/2013-11-28/mans-4m-bitcoin-fortune-lost-in-rubbish-dump.
- McCook, H. (2014) “An Order-of-Magnitude Estimate of the Relative Sustainability of the Bitcoin Network: A Critical Assessment of the Bitcoin Mining Industry, Gold Production Industry, the Legacy Banking system, and the Production of Physical Currency.” Available at www.bitcoin.fr/public/divers/docs/Estimation_de_la_durabilite_et_du_cout_du_reseau_Bitcoin.pdf.
- Merklejohn, S.; Pomarole, M.; Jordan, G.; Levechenko, K.; McCoy, D.; Voelker, G. M.; and Savage, S. (2013) “A Fistful of Bitcoins: Characterizing Payments among Men with No Names.” Paper prepared for IMC 13, October 23–25, Barcelona, Spain.
- Möser, M. (2013) “Anonymity of Bitcoin Transactions: An Analysis of Mixing Services.” Paper prepared for the Münster Bitcoin Conference, July 17–18.
- Nielsen, M. (2013) “How the Bitcoin Protocol Actually Works.” Blog posting at www.michaelnielsen.org/ddi/how-the-bitcoin-protocol-actually-works.
- Rivest, R. L.; Shamir, A.; and Adleman, L. A. (1978) “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems.” *Communications of the ACM* 21(2): 120–29.

White, L. H. (2014) "The Troubling Suppression of Competition from Alternative Monies: The Cases of the Liberty Dollar and E-gold." *Cato Journal* 34 (2): 281–301.

_____ (2015) "The Market for Cryptocurrencies." *Cato Journal* 35 (2): 383–402.